Do Credit Conditions Move House Prices?

Dan Greenwald* and Adam Guren†‡

March 7, 2019

Preliminary and Incomplete

Abstract

To what extent did an expansion and contraction of credit drive the 2000s housing boom and bust? The existing literature lacks consensus, with findings ranging from credit having no effect to credit driving the entire house price cycle. We show that the key difference behind these disparate results is how rental markets are modeled: assuming perfect segmentation between rental and owner-occupied housing leads to large effects of credit on house prices, while assuming frictionless rental markets makes credit irrelevant for house prices. We develop a model with frictional rental markets that nests both extremes and allows us to consider intermediate cases. We argue that the relative elasticity of the price-to-rent ratio and homeownership with respect to an identified credit shock is a sufficient statistic to measure these frictions. We estimate this moment, and use it to calibrate our model. Our results imply that rental markets are highly frictional and close to segmented, consistent with large effects of credit on house prices. Experiments using the structural model imply that credit conditions explain 47% - 57% of the rise in price-rent ratios over the boom.

*MIT Sloan School of Management

†Boston University

‡dlg@mit.edu, guren@bu.edu. We would like to thank Joe Vavra, Monika Piazzesi, Tim McQuade, and seminar participants at the 2019 AEA Annual Meeting for helpful comments. We would like to thank CBRE Economic Advisers and CoreLogic for providing and helping us understand their data. Liang Zhong provided excellent research assistance. Guren thanks the National Science Foundation (grant SES-1623801) for financial support.

1 Introduction

To what extent did an expansion and contraction of credit drive the 2000s housing boom and bust? This question is central to understanding the dramatic movements in housing markets that precipitated the Great Recession, yet a decade on there is no consensus on its answer. Some papers, such as Favilukis, Ludvigson, and Van Nieuwerburgh (2017), argue that changes in credit conditions can explain essentially all of the movements in house prices in the 2000s.[1] By contrast, Kaplan, Mitman, and Violante (2019) argue that credit conditions explain none of the boom and bust in house prices. Credit plays an important role in their model for the dynamics of homeownership, leverage, and foreclosures, but does not affect not house prices or the price-to-rent ratio, which are instead driven by beliefs.

The key difference behind these disparate findings is how these papers model the rental market. In models with frictionless rental markets, rental landlords step in to buy houses when credit contracts and sell houses when credit expands. If landlords' valuation of homes is insensitive to credit supply, the prices they offer for homes is unaffected by credit and depends only the present value of rents. When landlords are willing to transact an unlimited amount of housing at this credit-insensitive price, a credit-induced shift in demand for homes by constrained agents can move the homeownership rate (as these agents buy from landlords) but has essentially no effect on the equilibrium price. This can be thought of as a shift in borrower demand along a perfectly elastic landlord supply curve in price-to-rent ratio vs. homeownership rate space.

By contrast, if landlords cannot step in to buy houses when credit conditions change — because landlords have trouble converting owner-occupied homes to renter-occupied units or face credit constraints themselves — then shifts in credit supply lead to strong responses of house prices. Intuitively, the shift in housing demand by constrained households leads to a shift along an inelastic supply curve, leading to adjustment along the price margin. Most papers in the literature to date consider one of the two extremes: a perfectly inelastic supply curve or a perfectly elastic supply curve.

[1] Greenwald (2018), Guren, Krishnamurthy, and McQuade (2018), Garriga, Manuelli, and Peralta-Alva (2019), Garriga and Hedlund (2017), Garriga and Hedlund (2018), and Justiniano, Primiceri, and Tambalotti (2015a) also analyze models in which credit conditions play a key role in the boom and bust.

In this paper, we weigh in on this debate using using a new empirical moment and a model that allows us to nest both the perfectly inelastic and perfectly elastic extremes in addition to considering intermediate cases. In particular, we argue that the size of the causal effect of credit supply on the price-to-rent ratio relative to the causal effect of credit supply on the homeownership rate is a highly informative moment that disciplines the response of house prices to credit. This is because this ratio corresponds to the slope of a supply curve in price-to-rent ratio vs. homeownership rate space, which in turn determines the extent to which credit-induced shifts demand translate into movements in house prices. This slope therefore provides an important moment to be matched by any model seeking to study the influence of credit on house prices, regardless of the specific features of the model.

Rather than making modeling assumptions to pin down this slope, we treat it as an empirical object and provide new empirical estimates of the causal effect of credit supply on prices, rents, and homeownership rates. To this end, we create a panel of cities with data on prices, rents, homeownership rates, and credit supply. The rental measure is to our knowledge new to the macroeconomics literature, and improves on alternative series by measuring rents commanded by newly-rented units while excluding units on previously signed leases. To identify changes in credit, we use an instrument that uses differential city-level exposure to changes in the conforming loan limit interacted with changes in the national conforming loan limit following Loutskina and Strahan (2015). Our estimates indicate that shocks to credit supply increase prices, rents, and the price-to-rent ratio, but have no statistically significant effect on the homeownership rate. In terms of point estimates, we find that price-to-rent ratios respond at least five times as much as homeownership rates to a credit supply shock. We use this ratio of five-to-one as the key empirical moment to pin down the level of frictions in the rental market.

With these estimates in hand, we construct a dynamic equilibrium model, building on Greenwald (2018), in which house prices, rents, and the homeownership rate are all endogenous. Our primary modeling contribution is to tractably incorporate heterogeneity in landlord and borrower preferences for ownership, which allows our model to reproduce a fractional and time-varying homeownership rate. Our framework nests both perfect segmentation and frictionless conversion between rental and owner-occuped housing, as well as a continuum of intermediate cases.[2] We calibrate our model to match our key empirical moment, then use the model to compute the role of credit in driving the 2000s housing boom. We find that a relaxation of credit standards explains 47% to 57% of the rise in the price-to-rent ratio observed in the boom, with the precise number depending on our assumptions about other forces in the model such as changes in interest rates and house price expectations.

Our results imply we are in a world with significant segmentation between renter-occupied and owner-occupied housing. In fact, our Benchmark calibration generates house price dynamics that are much closer to those under the extreme of full segmentation than under the assumption of a frictionless rental market. At the same time, our Benchmark model generates a sizable and realistic movement in the homeownership rate and significant rent-to-own conversion in the boom and own-to-rent conversion in the bust, which is consistent with the data (see, e.g., Guren and McQuade (2018) and Kaplan et al. (2019)) but would be impossible under full segmentation. The ability of our model to jointly capture both price and homeownership dynamics is crucial to our results and implies an important advantage for "intermediate" models like ours relative to the polar cases typically observed in the literature.

The rest of the paper is structured as follows. Section 2 presents the supply and demand model diagrammatically in order to generate intuition and to motivate our estimation of the causal effects of credit on the homeownership rate and price-to-rent ratio. Section 3 describes the data, and Section 4 describes our instrument and empirical methodology. Section 5 presents our empirical results. Section 6 presents the model, Section 7 describes its calibration, and Section 8 presents our model results. Section 9 concludes.

[2]One could also generate intermediate cases through other mechanisms such as landlord credit or a strong intensive margin of housing consumption for savers. Indeed, Kaplan et al. (2019) have a strong saver intensive margin, which is why their finding that credit does not have a large effect on house prices is robust to alternate assumptions about rental markets. We argue that calibrating to our key empirical moment disciplines the model response regardless of which mechanism is at work, and we plan to evaluate these alternate mechanisms explicitly in future work.

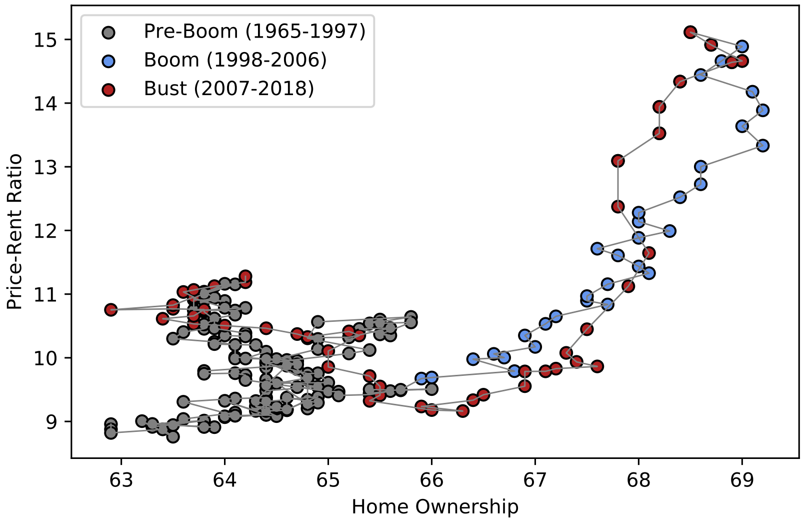

Figure 1: Price-Rent Ratio vs. Homeownership Rate

Note: The figure displays national data at the quarterly frequency. The price-rent ratio is obtained from the Flow of Funds, as the ratio of the value of housing services to the value of residential housing owned by households. The homeownership rate is obtained from the census (FRED code: RHORUSQ156N).

2 Intuition: Supply and Demand

Before we turn to the empirics and model, this section explains the intuition for how the rental market influences transmission from credit into house prices. This intuition motivates the structure of our model as well as our empirical focus on the causal effects of credit supply on the price-to-rent ratio and homeownership rate as the crucial sufficient statistics for calibration.

To begin, Figure 1 displays the evolution of the price-rent ratio and homeownership rate since 1965. Assuming that housing is either owned by households or by landlords/investors, each point on this plot represents an equilibrium between demand, that is the price the marginal renter is willing to pay to own a home, and supply, that is the price at which the marginal landlord is willing to sell that home.

The figure shows that these equilibria were fairly stable in the pre-boom era (1965 - 1997), with most observations clustered in the lower left portion of the figure. This pattern changed dramatically during the 1997 - 2006 housing boom. During this period, the price-rent ratio and homeownership rate increased in tandem to unprecedented levels. Following the start of the bust in 2007, these variables reverse course, traveling nearly the same path downward that they ascended during the boom, and finally ending up close to the typical values from the pre-boom era.

To understand what forces could have caused these patterns, we present a simple supply and demand treatment that illustrates the key forces in the equilibrium model we develop in Section 6. As in Figure 1, we use the price-rent ratio on the y-axis and the homeownership rate on the x-axis. The use of the price-rent ratio instead of house prices and the homeownership rate instead of quantities of owned housing ensures that changes are driven by the rent vs. own margin rather than the construction margin. Without this normalization, increases in construction might push up the quantity of owned housing while pushing down the price of housing, despite having no clear impact on the balance between owning and renting. Instead, this normalization allows us to focus the cost of owning *relative to* renting and the share of households who own — the key values for our margin of interest.

Demand for owner-occupied housing in this model comes from constrained households who require mortgages to own. As the price-to-rent ratio rises, more of these households prefer renting to owning, creating a downward slope. An expansion of credit supply shifts the curve rightward by allowing more favorable financing terms (either on interest rate or quantity of available credit) at a given price-to-rent ratio, inducing more households to choose ownership.

Supply in our model comes from landlords who decide whether to convert units of rental housing to owner-occuped housing, and sell it to households. The slope of the supply curve reflects the willingness of landlords to convert and sell more units as the price-to-rent ratio rises. The supply curve is shifted by anything that changes the landords' fundamental value of houses relative to rents. If landlords require credit, a credit supply shock would also shift the supply curve upward. Although potentially important, we abstract from landlord credit for the time being and return to it later.

Our supply and demand framework is displayed graphically in Figure 2. To begin, Figure 2a shows the case of perfect segmentation, in which units cannot be converted between owner-occupied and renter-occupied, and the homeownership rate is exogenously fixed. This example nests specifications such as Favilukis et al. (2017) and Justiniano, Primiceri, and Tambalotti (2015b), in which households cannot rent housing, corresponding to a fixed homeownership rate of 100%. In our framework, this corresponds to a perfectly inelastic supply, indicated by the vertical line in Figure 2a. This curve intersects the downward sloping demand curve to generate an equilibrium in

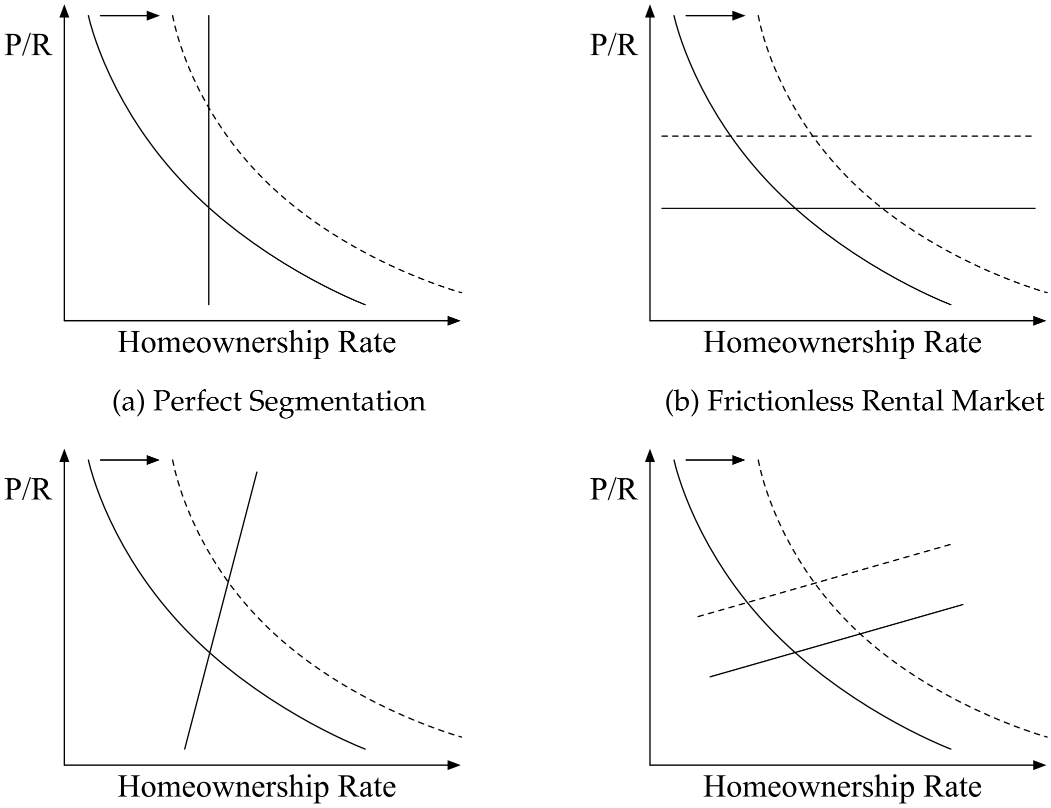

(a) Perfect Segmentation

(b) Frictionless Rental Market

(c) Intermediate Case 1

(d) Intermediate Case 2

Figure 2: Supply and Demand

price-rent vs. homeownership rate space.

From this starting point, we can consider the impact of a credit expansion. Assuming for now that only households use credit, the impact of this expansion is an outward shift in demand, as improved access to financing makes more households willing to purchase instead of renting at a given price. Under a perfectly inelastic (segmented) supply curve, this increased demand translates directly into an increase in house prices, while the homeownership rate remains fixed. Clearly, a credit expansion in this specification cannot reproduce the dual increases in both price-rent ratios and homeownership displayed in Figure 1.

Next, we can turn to Figure 2b, which corresponds to a frictionless rental market in which identical risk-neutral and deep-pocketed landlords transact with households, similar to the baseline model of Kaplan et al. (2019). This specification features a perfectly elastic (horizonal) supply curve, since these landlords are willing to buy or sell an unlimited amount of housing at a price equal to the present value of rents. Until the point at which landlords are completely driven out of the model and all housing is owner-occupied, house prices are pinned down by this present value relation. Since this present value does not depend on credit, an expansion of credit, corresponding to an outward shift of demand, increases the homeownership rate in this model, but cannot move the price-rent ratio.

Consequently, a credit expansion cannot explain the dual rise in price-rent ratios and homeownership under this specification either. Instead, reproducing the empirical pattern requires a *separate* upward shift in the supply curve, indicated by the horizontal dashed line in Figure 2b. Since prices are equal to the present value of rents in this model, what is required to move prices *relative* to current rents is to shock to expected future rents.[3] This analysis provides intuition for the finding in Kaplan et al. (2019) that a shock to housing beliefs (expected future rents) is responsible for the entire rise in price-rent ratios over the boom, while a credit expansion in their model affects the homeownership rate but not the price-rent ratio.

In this paper, we offer an intermediate approach that falls between the extreme specifications of perfect segmentation and frictionless rental markets. Specifically, we allow for *landlord heterogeneity*. While there are several potential sources for this heterogeneity, the most salient one is dispersion in the suitability of properties for renting: urban multifamily units can be easily rented at low cost, while detached or rural units may face much greater obstacles with respect to maintenance and moral hazard. This type of heterogeneity leads to an upward sloping supply curve such that the price landlords are willing to pay for housing varies with the identity of the marginal landlord or property. At high homeownership rates, the marginal converted property is easy to convert and maintain, and is valued highly by landlords relative to the rent it produces. At low homeownership rates, the marginal converted property is relatively costly to maintain as a rental property, and landlords are willing to part with it at a lower price-to-rent ratio.

The resulting supply-demand system is displayed in Figure 2c. With landlord heterogeneity, a credit expansion can simultaneously explain a rise in price-rent ratios and the homeownership rate without any second shock to supply, as the equilibrium travels along the upward sloping supply curve. Households buy up the properties that

[3]A shift in risk premia for the same set of rental cash flows would also generate a rise in present values if we had not assumed risk neutrality.

are least well-suited to renting, shifting the marginal property to one that is relatively more valued by landlords, increasing the marginal price-to-rent ratio. As a result, this model, unlike the previous specifications considered, can explain the empirical pattern observed during the housing boom with only a single shock.

At the same time, it is by no means clear that shocks to supply (e.g., to beliefs) did not play a major role during the boom. Indeed, any intermediate combination of upward sloping supply *alongside* a shock to supply can generate the same overall movement in price-rent and homeownership, as shown in Figure 2d.

What is needed to tell apart shifts in supply from movement along the supply curve is a measure of the slope of the supply curve. As is typical in the simultaneous equations literature, the slope of the supply curve can be uncovered using a shock to the demand curve. In Section 4 we propose a credit supply shock that provides exactly this type of variation and use it to estimate the elasticity of supply.

Before moving on, it is worth considering two extensions that we consider in detail later in the paper. First, consider the case where landlords are not deep-pocketed and require credit to buy and convert houses. In this case, a credit supply shock will shift the supply curve upward in addition to shifting the demand curve upward. Importantly, the credit supply shock we introduce in Section 4 is specific to borrower rather than landlord credit, so this modification should not distort our estimates of the supply slope. However, incorporating landlord use of credit would imply that credit conditions are even more important for house prices than is indicated by our estimates, and our results should be interpreted as a lower bound for the effect of credit on house prices.

Second, consider the case where there are unconstrained households ("savers") who value additional units of housing at a relatively constant marginal utility and are not credit constrained at the margin. In this case, even if rental markets are highly frictional, the housing supply curve may still be close to horizontal because savers are willing to absorb or supply housing to the constrained borrower households as credit supply fluctuates. An elastic intensive margin of housing consumption can thus flatten the supply curve. However, even if the slope of the housing supply curve is influenced by savers, this slope is *still* a sufficient statistic for the effect of credit supply expansions on the price-to-rent ratio and homeownership rate. Our approach of using this slope as an empirical moment is therefore informative regardless of whether the slope is determined by landlord or saver decisions, even though we do not directly feature this saver margin in our model.

In summary, the relative response of the price-to-rent ratio and the homeownership rate to an identified credit supply shock is the key moment we need to calibrate the slope of the supply curve, which is controlled by the degree of segmentation in rental markets. The remainder of the paper estimates this relative causal effect in an empirical analysis, while our structural model, which quantitatively implements our supply-demand framework, is calibrated to match our estimated supply curve slope to pin down the contribution of credit to the housing boom and bust.

3 Data

Our main data set is an annual panel at the core-based statistical area (CBSA) level. In order to examine the effect of credit supply on price to rent ratio and homeownership rate, the data set merges together data on house prices, rents, homeownership rates, and credit. We are limited by the availability of high-quality rent and homeownership rate data, and so our main sample covers 57 CBSAs over the period 1991-2016. Because coverage of homeownership and rents in our data changes over time, our main results are for an unbalanced panel. The details of data construction are in the Appendix.[4]

For house prices, we use the CoreLogic repeat sales house price index collapsed to an annual frequency. For rents, we use the CBRE Economic Advisors Torto-Wheaton Same-Store rent index (TW index). This is a high-quality rent index for multi-unit apartment buildings that is constructed similarly to an arithmetic repeat sales index.[5] It is available quarterly for 53 CBSAs beginning in 1989 and 62 CBSAs beginning in 1994. Importantly, the TW index differs in two ways from other indices. First, its re-

[4]In some cases data is only available for a broader metropolitan area rather than a smaller metropolitan division (e.g. Los Angeles and Orange County are divisions and they combine to make the Los Angeles Metropolitan Area). In these cases we use the data for the broader metro area for all of its constituent divisions.

[5]In particular, CBRE EA uses Axiometrics data on effective rents, which are asking rents for newly-rented units net of other leasting incentives. They build a historical rent series for each building and compute the index using changes in rents for identical units in the same buildings. Any missing quarters are imputed using the the index, as with a repeat sales index. CBRE does not use the standard repeat sales methodology because rent data is available for most buildings continuously; the imputation is only used for buildings that enter or leave the sample.

peat sales methodology is preferable to median rents, which are biased by changes in the composition of buildings rented out either due to conversion or construction. Second, the TW index uses asking rents on newly-leased apartments. This is important because rents are often fixed during the duration of a lease, and landlords often do not fully increase the rent when a tenant remains in a building. Because of this, measures of average or median rents such as the Bureau of Labor Statistics' median rents and the Zillow rent index are less volatile and tend to lag our rent measure. Because the price-to-rent ratio is meant to capture the rent a unit could command if it were leased out instead of sold, using newly-leased apartments is more appropriate. We compare the TW index to other rent measures in the Appendix. We also show that the TW index, which uses multi-unit apartment buildings, comoves strongly with a repeat-sales rent index for single family homes created by CoreLogic but available for a smaller subset of CBSAs and time periods, which suggests that the rental market is not highly segmented between single- and multi-family buildings and that using multi-family apartment rents is not a significant limitation.

Our homeownership data come from the Census' Housing and Vacancy Survey. The Census provides annual estimates of the homeownership rate at the CBSA and state levels from 1986 to 2017. A challenge that we must deal with is that the Census CBSA definitions change every decade, and there is no way to obtain homeownership rates for a consistent set of CBSA definitions. At times, the CBSAs expand substantially leading to large changes in homeownership rates due to the inclusion of additional counties. We deal with these issues in two ways. First, we do our best to harmonize the homeownership rates across different CBSA definitions. We do so by assuming that while the level of the homeownership rate may change, the annual log change in the homeownership rate remains the same. This means that we only need to adjust homeownership rates in years in which the definitions change, and we check that our results are robust to dropping these years in the Appendix. Second, we redo our analysis using a state-level panel in the Appendix. Because the state definitions do not change, the homeownership rates are not subject to these data issues. The downside to using state-level data is that we do not have a TW rent index at the state level.

For credit we use the Home Mortgage Disclose Act (HMDA) data, which we collapse to the CBSA levels. Following Favara and Imbs (2015), we use three different measure of credit: the dollar volume of loan originations, the number of loans, and the loan-to-income ratio. We use average income from the IRS rather than HMDA income as in Favara and Imbs. We also use HMDA for the construction of our instrument as described below.

4 Empirical Approach

4.1 Instrument

Our goal is to identify the effect of a local credit supply shock in city i at time t on prices, rents, price-to-rent ratios, and homeownership rates. To do so, we regress credit on these various housing market outcomes. Doing so using ordinary least squares (OLS) is problematic because credit is endogenous. In addition to credit affecting prices, rents, and homeownership rates, these outcomes likely affect credit supply. Furthermore, existing credit measures are noisy proxies for true credit availability. Loan volume and the number of loans are both affected by prepayment, which can be thought of credit replacement rather than a true expansion in credit. Consequently, we need an instrument for credit to credibly estimate the effect of credit on these outcomes in the face of endogeneity and measurement error.

Our empirical approach is to use a share-shift instrument that takes advantage of the fact that changes in the nation-wide Fannie Mae and Freddie Mac conforming loan limit (CLL) have bigger effects in cities with more homes priced near the limit following Loutskina and Strahan (2015). The government-sponsored enterprises (GSEs) Fannie Mae and Freddie Mac offer subsidized mortgage credit to mortgages that "conform" to requirements set by their regulator, the Federal Housing Finance Agency (FHFA, formerly OFHEO). Among the requirements for a loan to be conforming is that it must be below a maximum value, the CLL. Increases in the CLL induce changes in the supply of subsidized credit, expanding the overall supply of credit.

We follow Loutskina and Strahan in the construction of this instrument. To measure the mass of homes with loans near CLL, Loutskina and Strahan propose using the fraction of mortgage originations in the prior year that are within 5 percent of the current year's CLL in the HMDA data. Although this does fluctuate year-to-year, much of the variation is across cities. For instance, on average over our sample 7.2% of loans in San Francisco are originated within 5% of the next year's conforming loan limit. In El Paso, that figure is 0.4%. Our instrument is based on the idea that a change in the conforming loan limit should have a bigger effect in San Francisco than El Paso once one controls for time and city fixed effects.

In constructing the instrument, we have to make some changes relative to Loutskina and Strahan, who consider 1994-2006, in order to account for changes in the CLL implemented by Congress as part of the HERA legislation in 2008. In particular, Congress created more direct instructions for the national CLL and increased the CLL above the national CLL in certain high cost metropolitan areas. Future increases in the national CLL were tied to changes in a national house price index, although the CLL was not allowed to fall and was only allowed to rise when it passed its previous high watermark. However, high-cost cities could see their CLL rise by more than the national CLL up to a cap if their local house price index grew sufficiently quickly.[6] This would violate an instrumental variable's exclusion restriction because the change in the CLL would be mechanically correlated with lagged local outcomes. Consequently, in constructing the instrument we use the fraction of loans within 5 percent of the conforming loan limit using the actual limit in each city interacted with the change in the *national* CLL regardless of the change in the local CLL in high-cost areas. The fact that we are not using the actual change in the CLL in each city will reduce the instrument's power but maintain its exogeneity.[7]

We use two different but closely related empirical approaches. In the first, we follow the literature in regressing credit on contemporaneous outcomes instrumenting with contemporaneous shocks using a panel IV approach. However, many of the outcomes we consider such as house prices and price to rent ratios are notorious for sluggish adjustment and momentum (see, e.g., Guren (2018)). Our second specification is thus to plot the impulse response of the housing market outcomes we consider to credit using a panel local projection instrumental variables approach.

[6]The cap was also increased and then reduced in 2011.

[7]High-cost CBSAs have their CLL increase by the same percentage amount as the national CLL by default unless local house price growth triggers the local CLL calculation to be used. It turns out that in most high-cost areas in the period we study the national CLL change is used. Using the national CLL change in the few cities in which the local CLL change calculation is used does not significantly reduce the power of our instrument because there are relatively few such cities.

4.2 Panel Instrumental Variables for Contemporaneous Outcomes

For the first specification, our second-stage regression is:

$$\Delta \log(outcome_{i,t}) = \xi_i + \psi_t + \beta \Delta \log(credit_{i,t}) + \gamma \Delta \log(outcome_{i,t-1}) + \theta X_{i,t} + \epsilon_{i,t}, \quad (1)$$

where Δ represents changes between years t and $t-1$, $outcome_{i,t}$ is either a house price, rent, price-to-rent ratios or homeownership rate, $X_{i,t}$ is a vector of controls, and $credit_{i,t}$ is a measure of credit supply: either the number of originated loans, the volume of originated loans, or the average loan-to-income ratio. The CBSA fixed effect controls for average growth rates in the outcome across CBSAs. The time fixed effect picks up common national shocks. Including the lagged outcome growth as in Favara and Imbs (2015) controls for momentum in house prices and price to rent ratios and isolates the effect of credit rather than lagged shocks that are continuing to affect a CBSA.

The first stage is:

$$\Delta \log(credit_{i,t}) = \phi_i + \chi_t + \gamma Fraction_{i,T-1} \times \%\Delta CLL_t + \gamma \Delta \log(outcome_{i,t-1}) + \omega X_{i,t} + e_{i,t}. \quad (2)$$

While including time fixed effects absorbs the national change in the CLL, the fraction of loans in the previous year within 5% of this year's CLL is not accounted for by the regressors in (1). We do not want variation in this fraction alone to be part of the instrument and used for identification. Consequently, we include the fraction as a control variable in $X_{i,t}$ in both the first and second stage. We also include in $X_{i,t}$ the tripple interaction of the fraction, the national change in the CLL, and the local housing supply elasticity calculated by Saiz (2010). We do so because changes in credit supply may cause an increase in house prices that is stronger in cities with more inelastic housing supply, and the increase in house prices may mechanically lead to an expansion of credit on the intensive margin because the amount of loan one can obtain with a fixed loan-to-value ratio rises. We include this additional control to exclude this variation from the instrument. In order to make sure that our instrument represents the effect of expanding the conforming loan limit in an average-supply-elasticity city, we include the Saiz elasticity as a z-score (demeaned and divided by its standard error). Because these two variables appear both in the first and second stage, they act as control variables rather than instruments.

Our identifying assumption is that conditional on these controls there is no unobservable that varies with both the fraction of loans originated last year near the CLL and that varies with changes in the national CLL in the time series. If there were such an omitted variable, it would be picked up by our instrument, leading to biased results. In the Appendix, we conduct a number of robustness checks, including adding city-level income shocks and time-varying controls for the industrial structure.

As is typical for two-stage least squares with a single instrument, the IV coefficient of interest β is equal to the ratio of the coefficient on the instrument in a reduced-form regression of the instrument on the outcome to the first-stage coefficient γ. Under this interpretation, using different credit measures for the endogenous variable $\Delta \log(credit_{i,t})$ simply scales the reduced form effect of our instrument on the outcome into interpretable units of credit. Ultimately we are interested in the reduced form effect of our instrument on the outcome, and the rescaling is not crucial to our overall interpretation of the results.

Given the logic of the instrument, γ should be positive, as there should be a bigger effect of a national change in the conforming loan limit in cities in which more households are near the conforming loan limit. This turns out to be the case in practice for all three of our outcome measures. Table 1 shows the first-stage γ coefficient for each of our three credit measures for the price-to-rent ratio outcome.[8] In all cases, the first-stage coefficient and precision are fairly similar, which means the results will be reasonably close. We confirm this in regression tables and subsequently use the loan volume for our impulse response figures. Table 1 also shows the first-stage F statistics for the hypothesis that $\gamma = 0$. The F statistics are all between 11 and 14, suggesting that we do not face weak instrument issues (F less than 10).

4.3 Panel Local Projection Instrumental Variables

Our second specification is a panel local projection instrumental variables (LP-IV) approach. This approach generalizes the Jordà (2005) local projection methods to use exogenous instrumental variables for identification as in Ramey (2016) and Ramey

[8]The lagged outcome is included as a control variable in the first stage. The first stage looks similar for other outcomes so we do not report them.

Table 1: First Stage

Outcome	Number	Volume	LIR
$Fraction_{i,T-1} \times \%\Delta CLL_t$	46.379***	56.117***	55.262***
	(14.017)	(14.945)	(15.331)
F	10.949	14.099	12.993
N	1404	1404	1346

Notes: * Significant at 10% level, ** Significant at 5% level, *** Significant at 1% level. The table shows the coefficient on the instrument in the first stage regression (2) along with the first stage F statistic. The Y variable used (the lag of which is introduced as a control) is the price-to-rent ratio. All standard errors are robust.

and Zubairy (2018). Stock and Watson (2018) formalize the identification conditions for LP-IV. We extend this to the panel context and add CBSA and time fixed effects following Chen (2018).

In particular, for $k = 0, ..., 5$ we regress:

$$\log(outcome_{i,t+k}) = \xi_i + \psi_t + \beta_k \Delta \log(credit_{i,t}) + \theta X_{i,t} + \epsilon_{i,t}, \tag{3}$$

where $\Delta log(credit_{i,t})$ is the shock to credit in city i between $t-1$ and t. We instrument the credit shock using the Loutskina and Strahan (2015) instrument, so that the first stage is:

$$\Delta \log(credit_{i,t}) = \phi_i + \chi_t + \gamma Fraction_{i,T-1} \times \%\Delta CLL_t + \omega X_{i,t} + e_{i,t}. \tag{4}$$

We display our results by plotting the coefficients β_k along with their 95 percent confidence intervals against k. Because the first stage in the panel LP-IV approach is essentially identical to the first stage in the panel IV approach for contemporaneous outcomes, we do not report the first stage regressions.

The key identification conditions for the panel LP-IV specification following Stock and Watson (2018) are not only relevance and contemporaneous exogeneity but also exogeneity at all leads and lags. This requires that our instruments be independent of one another. Intuitively, serial correlation in the instrument would bias the estimated impulse response because the outcome in period $t+k$ would be affected by shocks to credit in periods other than time t. To account for this we not only include $Fraction_{i,T-1}$ and $Fraction_{i,T-1} \times \%\Delta CLL_t \times Saiz_i$ as controls in $X_{i,t}$ but we also include a lag of these two variables and our instrument as controls. This follows Ramey (2016) and Ramey and Zubairy (2018). Furthermore, to flexibly control for serial correlation and momentum in the outcome variable, we include two lags of the outcome, $\log(outcome_{i,t-1})$ and $\log(outcome_{i,t-2})$ as controls in $X_{i,t}$. Using these as controls helps ensure that the lead and lag exogeneity conditions are satisfied.

Nonetheless, even with these controls, lead exogeneity may be violated as changes in the CLL are permanent. As we mention above, the interpretation of our results as instrumental variables for credit is not crucial: the reduced-form effect of the permanent change in the CLL on the housing market is what is of interest to us and the IV procedure only converts the results into interpretable units of credit. Indeed, in calibrating our model we introduce a permanent change in the CLL in our model. Lag exogeneity violations, however, would bias our results, which is why we include such extensive controls for lags of the outcome and instruments.

4.4 Future Work: Additional Instruments

Most of the variation in our instrument comes from the "boom" period and not from the contraction of credit in the bust because the conforming loan limit rises when house prices rise nationally, does not fall when house prices fall, and only rises again when prices pass their high watermark. In future work, we plan to use additional instruments for credit supply that take advantage of the bust to make sure that our results generalize to the bust.[9]

5 Empirical Results

5.1 Panel IV Results

Table 2 shows OLS estimates of equation (1) for three credit measures (rows), the number of loans, the dollar volume of loans, and the loan-to-income ratio, and for four different outcomes (columns), the price-to-rent ratio (CoreLogic prices and Torto-Wheaton rents), the homeownership rate (Census), house prices, and rents. Table 3

[9] We have evaluated several candidate instruments. The main difficulty is that we only have a handful of cities. Many instruments in the literature vary systematically across states or regions and require state-by-time fixed effects so that the variation is limited to within-states. We do not have rents or homeownership rates for enough cities to have statistical power with state-by-time fixed effects.

Table 2: OLS Results For Contemporaneous Outcomes

	$\Delta\log$(Price/Rent)			$\Delta\log$(Homeownership Rate)		
$\Delta\log$(# Loans)	-0.004			-0.007		
	(0.007)			(0.004)		
$\Delta\log$(Vol. Loans)		0.020***			-0.006	
		(0.006)			(0.005)	
$\Delta\log$(Loan/Income)			0.018***			-0.003
			(0.006)			(0.005)
N	1404	1404	1346	1729	1729	1653
	$\Delta\log$(Price)			$\Delta\log$(Rent)		
$\Delta\log$(# Loans)	0.013**			0.014***		
	(0.005)			(0.005)		
$\Delta\log$(Vol. Loans)		0.038***			0.021***	
		(0.005)			(0.005)	
$\Delta\log$(Loan/Income)			0.028***			0.013**
			(0.005)			(0.005)
N	1404	1404	1346	1404	1404	1346

Notes: * Significant at 10% level, ** Significant at 5% level, *** Significant at 1% level. The Table shows ordinary least squares estimates of equation (1). The control variables $X_{i,t}$ include $Fraction_{i,t-1}$ and $Fraction_{i,Tt-1} \times \%\Delta CLL_t \times Z(Saiz_i)$. All standard errors are robust.

shows the corresponding IV estimates. The coefficients can be interpreted as the elasticity of the contemporaneous outcome variable to an increase in the credit measure.

The OLS results show that there are small but positive elasticities of the price-to-rent ratio to credit and a precise zero effect on homeownership rates. The price-to-rent ratio response results from a combination of a price response that is roughly double the price-to-rent ratio response together with a rent response that is commensurate to the price-to-rent ratio response.

The IV responses are qualitatively similar but quantitatively much larger. The fact that OLS is biased downward suggests that there is significant measurement error in credit that is being corrected by our instrument. This makes sense: there is significant prepayment and refinancing that does not represent a true expansion in credit.

In particular, our results show an elasticity of the price-to-rent ratio of between 0.12 and 0.15 depending on the credit measure, which is significant at the 10% level. Conversely, the elasticity of the homeownership rate to credit supply is 0.03 and insignif-

Table 3: IV Results For Contemporaneous Outcomes

	$\Delta\log(\text{Price/Rent})$			$\Delta\log(\text{Homeownership Rate})$		
$\Delta\log(\text{\# Loans})$	0.152*			0.032		
	(0.086)			(0.068)		
$\Delta\log(\text{Vol. Loans})$		0.116*			0.027	
		(0.071)			(0.058)	
$\Delta\log(\text{Loan/Income})$			0.142*			0.027
			(0.075)			(0.065)
N	1404	1404	1346	1729	1729	1653
	$\Delta\log(\text{Price})$			$\Delta\log(\text{Rent})$		
$\Delta\log(\text{\# Loans})$	0.378**			0.261***		
	(0.104)			(0.087)		
$\Delta\log(\text{Vol. Loans})$		0.302***			0.214***	
		(0.065)			(0.063)	
$\Delta\log(\text{Loan/Income})$			0.359**			0.254***
			(0.091)			(0.086)
N	1404	1404	6940	1404	1404	1346

Notes: * Significant at 10% level, ** Significant at 5% level, *** Significant at 1% level. The Table shows instrumental variables estimates of equation (1) where the instrument is $Fraction_{i,t-1} \times \%\Delta CLL_t$. The control variables $X_{i,t}$ include $Fraction_{i,t-1}$ and $Fraction_{i,t-1} \times \%\Delta CLL_t \times Z(Saiz_i)$. All standard errors are robust.

icant.[10] Using our point estimates, one can infer that the response of the price-to-rent ratio to credit is about four to five times as large as the response of homeownership rates to credit. The response of the price-to-rent ratio is driven by an elasticity of prices to credit of between 0.30 and 0.38 and a response of rents to credit of between 0.21 and 0.26. [11]

Our results for rents are large relative to the literature, which tends to find that rents are sticky and move relatively little. This is likely because of the rent measure we use. The Torto-Wheaton Rent Index measures the rent commanded by a newly-rented unit, while most rent measures include ongoing rental contracts in which the rent is fixed (typically for a year) and only partially renegotiated while existing ten-

[10] One concern mentioned in Section 3 is that the raw Census homeownership rate data uses changing CBSA definitions and the imperfect correction for these changing definitions we use may bias our results. Consequently, in the Appendix we show we find similar results using a state-level panel that is not subject to these data issues.

[11] One cannot simply subtract the coefficient for rents from the coefficient for prices to obtain the coefficient on the price to rent ratio because of the way the controls and fixed effects differentially effect each outcome variables. Nonetheless, the price-to-rent ratio results are to a first order close to the difference between the price and rent coefficients.

ants stay in place due to the costs of finding a new tenant. Our results suggest that rents for newly-rented units respond significantly to a credit supply shock. This result highlights the importance of using the right type of rent measures for constructing price-to-rent ratios in empirical work.

The magnitude we find for the price-to-rent ratio is consistent with what Favara and Imbs (2015) find for the response of house prices, although our estimated response of house prices is larger than what they find. Our results are different from what Gete and Reher (2018) find in the bust. Using the share of local loans originated by banks undergoing stress tests in the bust as an instrument for the loan denial rate, Gete and Reher find that loan denials cause an increase in rents, a decrease in homeownership rates, and a decrease in house prices. Our results may differ for three reasons. First, Gete and Reher use the denial rate, while we look at credit quantities. Second, Gete and Reher use Zillow data on average rents rather than rents for newly-rented units as in the Torto-Wheaton index we use. Third, while we focus on an instrument with significant variation in the boom, while Gete and Reher use an instrument for credit supply in the bust. In future work, we plan to use instruments that have variation in the bust.

5.2 Panel Local Projection IV Impulse Responses

Figure 3 shows the impulse responses using panel local projection instrumental variables.[12] The figures show results for the dollar volume of new loans as our measure of credit; the other credit measures look similar. The price-to-rent ratio exhibits a hump-shaped response to credit shocks, peaking at about three years at an elasticity of 0.47. This shows that the 0.12 elasticity we found over one year using the panel IV approach only shows the tip of the iceberg for the overall impulse response.[13] The price to rent ratio combines a hump-shaped response of prices that peaks at around 0.79 after two years and a jump in rents of about 0.2 on impact that stays persistent. The fact that house prices display momentum but rents do not suggests that rents for newly-rented units are not very sticky.

[12] As with the panel IV, the OLS impulse responses are an order of magnitude smaller than IV but qualitatively similar in that the price-to-rent ratio and rents respond similarly, prices respond by about twice as much, and homeownership rates show no significant response.

[13] The standard errors are somewhat wider for the LP-IV approach owing to the additional controls. However, the magnitudes we find over one year are roughly consistent with Table 3.

Figure 3: Panel Local Projection Instrumental Variable Impulse Responses

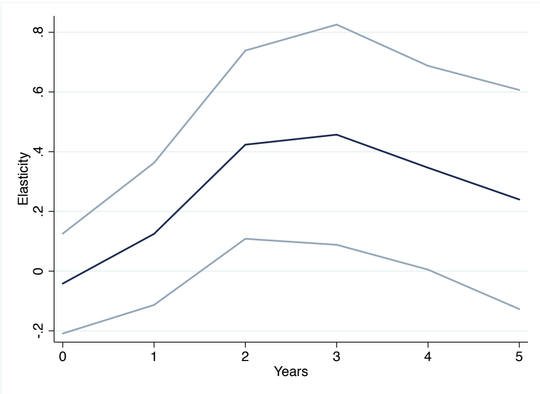

(a) Price/Rent

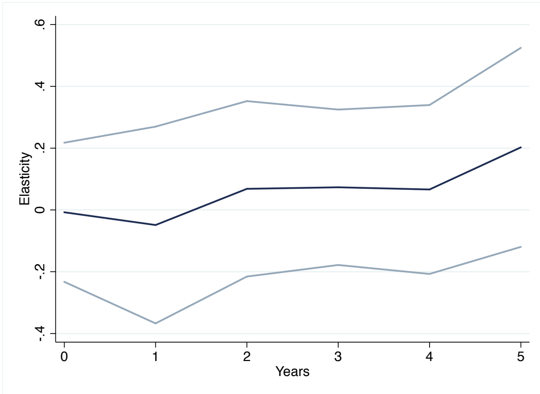

(b) Homeownership Rate

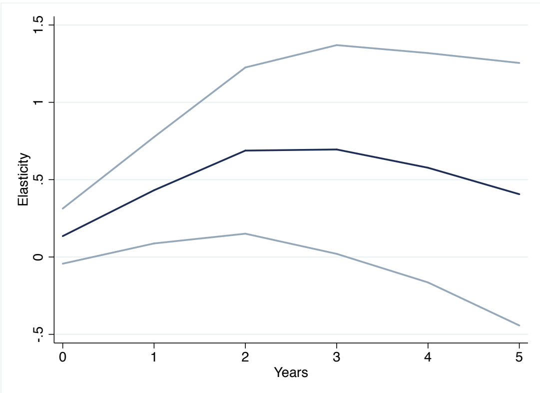

(c) House Price

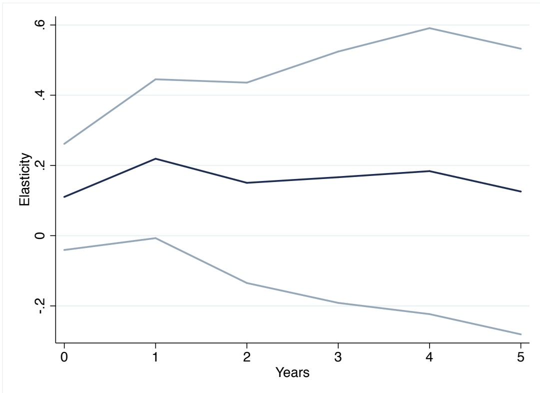

(d) Rent

Notes: 95 % confidence interval shown by gray bands. The figure shows panel local projection instrumental variables estimates of the response of the indicated outcomes to dollar credit volume. The second stage is as indicated in equation (3) and the first stage is as indicated in equation (4). The instrument is $Fraction_{i,t-1} \times \%\Delta CLL_t$ and control variables include $Fraction_{i,t-1}$, $Fraction_{i,t-1} \times \%\Delta CLL_t \times Z(Saiz_i)$, $Fraction_{i,t-2}$, $Fraction_{i,t-2} \times \%\Delta CLL_t - 1$, $Fraction_{i,t-2} \times \%\Delta CLL_t - 1 \times Z(Saiz_i)$, and two lags of the outcome variable. All standard errors are robust.

By contrast, we find no significant response of homeownership rates to credit shocks. Generously interpreting the point estimate for homeownership rates over three to four years gives a response that is about one fifth of that of the price-to-rent ratio, as with Table 3. We use this figure, a five-times larger response of the price-to-rent ratio than the homeownership rate, in our model calibration.

6 Model

This section develops an equilibrium model that we use to quantitatively evaluate the role of credit in driving house prices, with a particular focus on the recent boom-bust cycle. The model is designed so that it can capture our empirical findings on the relative elasticity of the price-to-rent ratio and the homeownership rate to credit shocks that shift household demand for housing.

Demographics. There is a representative borrower, landlord, and saver. Each are infinitely lived, permanent types, with perfect risk sharing within the members of each type.

Housing Technology. Housing is produced by construction firms (described subsequently) whose supply at the end of period t is denoted $\bar{H}_t$. Housing can be owned either by borrowers or by landlords, who in turn rent the housing they own to borrowers. We denote borrower-owned housing as $H_{B,t}$ and landlord-owned rental housing as $H_{L,t}$. Housing produces a service flow proportional to the stock, and is sold ex-dividend (i.e., after the service flow is consumed).

Preferences. Borrowers have log preferences over a CES aggregator of nondurable consumption and housing services:

$$U_B = \sum_{t=0}^{\infty} \beta_B^t \log \left(\left[\hat{c}_{B,t}^{\frac{\varepsilon-1}{\varepsilon}} + \xi \hat{h}_{B,t}^{\frac{\varepsilon-1}{\varepsilon}} \right]^{\frac{\varepsilon}{\varepsilon-1}} \right)$$

where c represents nondurable consumption, h represents housing services, and hats indicate that variables have been put in per-capita terms by dividing through by the population share χ_B for $j \in \{B, S\}$. Savers also are assumed to have log preferences, but do not consume housing:

$$U_S = \sum_{t=0}^{\infty} \beta_S^t \log \left(\hat{c}_{S,t} \right).$$

This restrictive and important assumption shuts down any margin for borrowers and savers to transact housing, equivalent to fully segmented housing markets between these two groups. The implications of this choice are discussed in detail in Section 7. For landlords, we use risk neutral preferences:

$$U_L = \sum_{t=0}^{\infty} \beta_L^t \hat{c}_{L,t}$$

following the baseline model of Kaplan et al. (2019) to allow easier comparison to their framework.

Asset Technology. Borrowers and landlords can trade long-term mortgage debt with savers at equilibrium, with the mortgage technology following Greenwald (2018). Borrower debt is denoted $M_{B,t}$ while landlord debt is denoted $M_{L,t}$. Debt is issued in the form of fixed-rate nominal perpetuities with coupons that geometrically decay at rate ν. This means that a mortgage that is issued with balance M^* and rate r^* will have payment stream of $(r^* + \nu)M^*,\ (1-\nu)(r^* + \nu)M^*,\ (1-\nu)^2(r^* + \nu)M^*, \ldots$. Mortgage loans are prepayable, with fraction $\rho_{j,t}$ of agents of type j choosing to prepay their loans in a given period. As in Greenwald (2018), the average size of new loans for agents i of type j (denoted $M^*_{i,j,t}$) is subject to both loan-to-value (LTV) and payment-to-income (PTI) limits at origination, of the form:

$$
\begin{aligned}
M^*_{i,j,t} &\leq \theta^{LTV}_{j,t} p_t H^*_{i,j,t} \\
M^*_{i,j,t} &\leq \frac{\left(\theta^{PTI}_{j,t} - \omega_j\right) \text{income}_{i,j,t}}{r^*_{j,t} + \nu_j + \alpha_j},
\end{aligned}
$$

where p_t is the price of housing, $H^*_{i,j,t}$ is the borrower's new house size, and ω_j and α_j are offsets used in calibration to account for non-housing debts, and taxes and insurance, respectively.

In addition to the mortgage contract, there is a one-period bond B_t in zero net supply. Agents cannot take a short position in this bond, so it is traded by the savers only at equilibrium. All financial contracts are nominal, meaning that real payoffs from both one-period bonds and mortgages decay each period at the constant rate of inflation π.

Finally, there is a divisible housing good, whose holdings by the borrower and landlord are denoted H_B and H_L, respectively. Only borrowers who are currently prepaying their existing loans are eligible to purchase new housing (i.e., borrowers face a constant probability of receiving a moving shock). This good produces a flow of housing services equal to its stock, and requires a per-period maintenance cost of fraction δ of the current value of housing.

Ownership Benefit Heterogeneity. Without additional heterogeneity, the model would be unable to generate a fractional and time-varying homeownership rate. Essentially, if all borrowers have the same valuation for housing, and all landlords have the same valuation for housing, then whichever group has the higher valuation will own all the housing, leading to a homeownership rate of either 0% or 100%. In order to generate a fractional homeownership rate, we need to impose further heterogeneity in how agents value housing *within* at least one of these types. Our key modeling contribution in this paper is to explicitly allow for this within-type heterogeneity.

We impose this heterogeneity in a simple way, by assuming that agents receive an additional service flow (either positive or negative) from owning housing. For parsimony, we assume that if borrower i owns one unit of housing, he or she receives surplus equivalent to $\omega_{i,t}$ of the numeraire, where $\omega_{i,t}^{B} \sim \Gamma_{\omega,B}$ is drawn i.i.d. across borrowers and time. Symmetrically, if landlord i owns one unit of housing, he or she receives surplus equivalent to $\omega_{i,t}^{L} \sim \Gamma_{\omega,L}$ of the numeraire. Since we perceive these costs as likely non-financial, we rebate them lump-sum to households, making them equivalent to utility benefits or costs.[14]

There are several forms of heterogeneity that would map intuitively into this framework. On the borrower side, heterogeneity in the value of ownership likely stands in for household age, family composition, ability to make a down payment, and true personal preference for ownership. On the landlord side, we suspect that the biggest source of heterogeneity is on the suitability of different properties for rental. For example, while urban multifamily units can be efficiently monitored and maintained in a rental state, the depreciation and moral hazard concerns for renting a detached suburban or rural house may be much higher.

The degree of dispersion of the distributions $\Gamma_{\omega,B}$ and $\Gamma_{\omega,L}$ map into the slopes of the demand and supply curves, respectively, in Section 2. Specifically, the more dispersed are the ownership benefits, the steeper the slope, as the marginal valuation changes rapidly as we move along the distribution. In contrast, a distribution with low dispersion will yield a flatter, more elastic curve, as agents share highly similar valuations, implying that prices move little as the homeownership rate, and the identities

[14]For timing, we assume that agents must pick out their target house size prior to drawing $\omega_{i,t}$, but can then choose whether to proceed with the transaction or back out. This ensures that a single landlord with the highest benefit cannot buy up the entire stock of rental housing, undoing the effects of heterogeneity at equilibrium.

of the marginal owner/renter and landlord vary.

Borrower's Problem. The borrower's budget constraint is:

$$
\begin{aligned}
c_{B,t} \leq\ & \underbrace{(1-\tau)y_{B,t}}_{\text{after-tax income}} + \underbrace{\rho_{B,t}\left(M^*_{B,t} - \pi^{-1}(1-\nu_B)M_{B,t-1}\right)}_{\text{net mortgage iss.}} - \underbrace{\pi^{-1}(1-\tau)X_{B,t-1}}_{\text{interest payment}} - \underbrace{\nu_B\pi^{-1}M_{B,t-1}}_{\text{principal payment}} \\
& - \underbrace{\rho_{B,t}p_t\left(H^*_{B,t} - H_{B,t-1}\right)}_{\text{net housing purchases}} - \underbrace{\delta p_t H_{B,t-1}}_{\text{maintenance}} - \underbrace{q_t\left(h_{B,t} - H_{B,t-1}\right)}_{\text{rent}} + \underbrace{\left(\int_{\bar{\omega}_{B,t-1}} \omega \, d\Gamma_{\omega,B}\right)\bar{H}_{t-1}}_{\text{owner surplus}} \\
& - \underbrace{\left(\int^{\Gamma_\kappa^{-1}(\rho_{B,t})} \kappa \, d\Gamma_\kappa(\kappa)) - \bar{\Psi}_t\right) M^*_{B,t}}_{\text{refi cost (rebated)}} + \underbrace{T_{B,t}}_{\text{other rebates}}
\end{aligned}
$$

where $y_{B,t}$ is exogenous outside income and q_t is the rental rate (i.e., the price of housing services). The optimal policy is for all borrowers with owner utility shock $\omega_{i,t} > \bar{\omega}_t$ to choose to buy housing. By market clearing, we have $\bar{\omega}_{B,t} = \Gamma^{-1}_{\omega,B}(1 - H_{B,t}/\bar{H}_t)$, which ensures that the fraction of borrowers choosing to own is equal to the fraction of borrower-owned housing. Income is taxed at rate τ, while interest payments on the mortgage are tax deductible. For refinancing, borrowers draw heterogeneous costs $\kappa_{i,t}$ so that the borrowers with the lowest draws refinance, following Greenwald (2018).

The laws of motion for mortgage balance $M_{B,t}$, interest payment $X_{B,t}$, and owned housing $H_{B,t}$ are:

$$
\begin{aligned}
M_{B,t} &= \underbrace{\rho_{B,t}M^*_{B,t}}_{\text{new loans}} + \underbrace{(1-\rho_{B,t})(1-\nu_B)\pi^{-1}M_{B,t-1}}_{\text{old loans}} \\
X_{B,t} &= \underbrace{\rho_{B,t}r^*_{B,t}M^*_{B,t}}_{\text{new loans}} + \underbrace{(1-\rho_{B,t})(1-\nu_B)\pi^{-1}X_{B,t-1}}_{\text{old loans}} \\
H_{B,t} &= \underbrace{\rho_{B,t}H^*_{B,t}}_{\text{new housing}} + \underbrace{(1-\rho_{B,t})H_{B,t-1}}_{\text{old housing}}.
\end{aligned}
$$

The borrower's optimality conditions are:

$$
(h_{B,t}): \qquad q_t = (u^h_{B,t}/u^c_{B,t}) \tag{5}
$$

$$
(H^*_{B,t}): \qquad p_t = \frac{\mathbb{E}_t\left\{\Lambda_{B,t+1}\left[\bar{\omega}_{B,t} + q_{t+1} + \left(1-\delta-(1-\rho_{B,t+1})\mathcal{C}_{B,t+1}\right)p_{t+1}\right]\right\}}{1-\mathcal{C}_{B,t}} \tag{6}
$$

$$(M^*_{B,t}): \qquad 1 = \Omega^B_{M,t} + r^*_{j,t}\Omega^B_{X,t} + \mu_{B,t} \tag{7}$$

$$(\rho_{B,t}): \qquad \rho_{B,t} = \Gamma_\kappa \Bigg\{ \underbrace{\left(1 - \Omega^B_{M,t} - \Omega^B_{X,t}\bar{r}_{B,t-1}\right)\left(1 - \frac{(1-\nu)\pi_t^{-1} M_{B,t-1}}{M^*_{B,t}}\right)}_{\text{new debt incentive}} - \underbrace{\Omega^B_{X,t}\left(r^*_{B,t} - \bar{r}_{B,t-1}\right)}_{\text{interest rate incentive}} \Bigg\}, \tag{8}$$

where

$$\mathcal{C}_{B,t} = \mu_{B,t} F^{LTV}_{B,t} \theta^{LTV}_{B,t}.$$

$\mu_{B,t}$ is the multiplier on the borrowing constraint, $F^{LTV}_{B,t}$ is the fraction of borrowers who are LTV-constrained, $\bar{r}_{B,t-1} = X_{B,t-1}/M_{B,t-1}$ is the average rate on existing debt, and the marginal continuation cost of principal balance $\Omega^B_{M,t}$ and of interest payments $\Omega^B_{X,t}$ satisfy:

$$\Omega^B_{M,t} = \mathbb{E}_t \left\{ \Lambda_{B,t+1}\pi^{-1}\Big[\nu_B + (1-\nu_B)\Big(\rho_{B,t+1} + (1-\rho_{B,t+1})\Omega^B_{M,t+1}\Big)\Big] \right\}$$
$$\Omega^B_{X,t} = \mathbb{E}_t \left\{ \Lambda_{B,t+1}\pi^{-1}\Big[(1-\tau) + (1-\nu_B)(1-\rho_{B,t+1})\Omega^B_{X,t+1}\Big] \right\}.$$

Equation (5) sets the rent equal to the marginal rate of substitution between housing services and consumption. Equation (6) specifies that at an interior solution, the price of housing must be equal to the present value of next period's service flow (the rent combined with the owner's utility bonus) plus the continuation value. Note that since p_t is the price of newly purchased housing that is about to be borrowed against, p_t includes the value of collateral services, which the borrower does not receive in periods when she does not refinance. Therefore, the continuation value is equal to the market value of housing net of maintenance costs, $(1-\delta)p_{t+1}$, minus the value of collateral services $\mu_{B,t+1}\theta_B$ in states of the world when the borrower does not refinance, which occurs with probability $(1-\rho_{B,t+1})$. Equation (7) sets the marginal benefit of one unit of face value debt ($1 today) against the marginal cost (the continuation cost of the debt plus the shadow cost of tightening the borrowing constraint). Equation (8) sets the transaction cost for the marginal refinancer equal to the net marginal benefit of refinancing, making this borrower indifferent.

Landlord's Problem. The landlord's problem is similar, with the key exception that the landlord only sells housing services to the borrower instead of consuming them.

The landlord's budget constraint is:

$$
\begin{aligned}
c_{L,t} \leq\ & \underbrace{(1-\tau)y_{L,t}}_{\text{after-tax income}} + \underbrace{\rho_{L,t}\left(M^*_{L,t} - \pi^{-1}(1-\nu_L)M_{L,t-1}\right)}_{\text{net mortgage iss.}} - \underbrace{\pi^{-1}(1-\tau)X_{L,t-1}}_{\text{interest payment}} - \underbrace{\nu_L \pi^{-1} M_{L,t-1}}_{\text{principal payment}} \\
& - \underbrace{\rho_{L,t} p_t \left(H^*_{L,t} - H_{L,t-1}\right)}_{\text{net housing purchases}} - \underbrace{\delta p_t H_{L,t-1}}_{\text{maintenance}} + \underbrace{q_t H_{L,t-1}}_{\text{rent}} + \underbrace{\left(\int_{\bar{\omega}_{L,t-1}} \omega \, d\Gamma_{\omega,L}\right) \bar{H}_{t-1}}_{\text{owner surplus}} \\
& - \underbrace{\left(\int^{\Gamma_\kappa^{-1}(\rho_{L,t})} \kappa \, d\Gamma_\kappa(\kappa)) - \bar{\Psi}_t\right) M^*_{L,t}}_{\text{refi cost (rebated)}} + \underbrace{T_{L,t}}_{\text{other rebates}},
\end{aligned}
$$

The landlord's optimality conditions are

$$
\begin{aligned}
(H^*_{B,t}): \quad & p_t = \frac{\mathbb{E}_t \left\{ \Lambda_{L,t+1} \left[\bar{\omega}_{L,t} + q_{t+1} + \left(1 - \delta - (1-\rho_{L,t+1})\mathcal{C}_{L,t+1}\right) p_{t+1} \right] \right\}}{1 - \mathcal{C}_{L,t}} \\
(M^*_{L,t}): \quad & 1 = \Omega^L_{M,t} + r^*_{j,t}\Omega^L_{X,t} + \mu_{L,t} \\
(\rho_{L,t}): \quad & \rho_{L,t} = \Gamma_\kappa \left\{ (1 - \Omega^L_{M,t} - \Omega^L_{X,t}\bar{r}_{L,t-1}) \left(1 - \frac{(1-\nu)\pi_t^{-1} M_{L,t-1}}{M^*_{L,t}}\right) - \Omega^L_{X,t}\left(r^*_{L,t} - \bar{r}_{L,t-1}\right) \right\},
\end{aligned}
$$

where $\mathcal{C}_{L,t} = \mu_{L,t} F^{LTV}_{L,t} \theta^{LTV}_{L,t}$ is defined analogously to the borrower case. The landlord's problem is also subject to the laws of motion:

$$
\begin{aligned}
M_{L,t} &= \rho_{L,t} M^*_{L,t} + (1-\rho_{L,t})(1-\nu_L)\pi^{-1} M_{L,t-1} \\
X_{L,t} &= \rho_{L,t} r^*_{L,t} M^*_{L,t} + (1-\rho_{L,t})(1-\nu_L)\pi^{-1} X_{L,t-1} \\
H_{L,t} &= \rho_{L,t} H^*_{L,t} + (1-\rho_{L,t}) H_{L,t-1},
\end{aligned}
$$

and the fixed point condition that pins down the marginal continuation cost of debt:

$$
\begin{aligned}
\Omega^L_{M,t} &= \mathbb{E}_t \left\{ \Lambda_{L,t+1} \pi^{-1} \left[\nu_L + (1-\nu_L)\left(\rho_{L,t+1} + (1-\rho_{L,t+1})\Omega^L_{M,t+1}\right) \right] \right\} \\
\Omega^L_{X,t} &= \mathbb{E}_t \left\{ \Lambda_{L,t+1} \pi^{-1} \left[(1-\tau) + (1-\nu_L)(1-\rho_{L,t+1})\Omega^L_{X,t+1} \right] \right\}.
\end{aligned}
$$

Saver's Problem. The saver's budget constraint is:

$$
c_{S,t} \leq \underbrace{(1-\tau)y_{S,t}}_{\text{after-tax income}} - \underbrace{(B_t - R_{t-1}B_{t-1})}_{\text{net bond purchases}} + \underbrace{T_{S,t}}_{\text{rebates}}
$$

$$+\sum_{j\in\{B,L\}}\left\{\underbrace{\pi^{-1}(\bar{r}_j+\nu_j)M_{j,t-1}}_{\text{total payment}}-\underbrace{\rho_{j,t}\left(\exp(\Delta_{j,t})M^*_{j,t}-\pi^{-1}(1-\nu_j)M_{j,t-1}\right)}_{\text{net mortgage iss.}}\right\}$$

where the wedge $\Delta_{j,t}$ is a time-varying tax, rebated to the saver lump sum at equilibrium, that allows for time variation in mortgage spreads. A value of $\Delta_{j,t} > 0$ implies that mortgage rates exceed the rates on risk-free bonds implied by the expectations hypothesis (or in other words, that mortgage bonds trade at a discount).

The saver's optimality conditions are:

$$\begin{aligned}(B_t):\quad & 1 = R_t\mathbb{E}_t\left[\pi^{-1}\Lambda_{S,t+1}\right]\\ (M^*_{j,t}):\quad & 1 = Q^S_{M,t} + r^*_{j,t}Q^S_{X,t}\end{aligned}$$

where the marginal continuation values of principal balance and promised interest payments are given by:

$$\begin{aligned}Q^S_{M,t} &= \mathbb{E}_t\left\{\Lambda_{S,t+1}\pi^{-1}\left[\nu_B+(1-\nu_B)\left(\rho_{B,t+1}+(1-\rho_{B,t+1})Q^S_{M,t+1}\right)\right]\right\}\\ Q^S_{X,t} &= \mathbb{E}_t\left\{\Lambda_{S,t+1}\pi^{-1}\left[(1-\tau)+(1-\nu_B)(1-\rho_{B,t+1})Q^S_{X,t+1}\right]\right\}.\end{aligned}$$

Construction Firm's Problem. New housing is produced by a continuum of competitive construction firms. Similar to Favilukis et al. (2017) and Kaplan et al. (2019), we assume that new housing is produced using nondurable goods Z and land L_t according to the technology:

$$\begin{aligned}I_t &= L_t^{\varphi}Z_t^{1-\varphi}\\ \bar{H}_t &= (1-\delta)\bar{H}_{t-1}+I_t\end{aligned}$$

where $\bar{L}$ units of land permits are auctioned off by the government each period, with the proceeds returned pro-rata to the households. Each construction firm therefore solves:

$$\max_{L_t,Z_t} p_tL_t^{\varphi}Z_t^{1-\varphi} - p_{L,t}L_t - Z_t$$

where $p_{L,t}$ is the equilibrium price of land permits, and the price of nondurables is normalized to unity. The construction firm's optimality conditions are therefore:

$$
\begin{aligned}
p_{L,t} &= p_t \varphi L_t^{\varphi-1} Z_t^{1-\varphi} \\
1 &= p_t (1-\varphi) L_t^{\varphi} Z_t^{-\varphi}.
\end{aligned}
$$

Equilibrium. A competitive equilibrium economy consists of endogenous states $(H_{B,t-1}, M_{B,t-1}, X_{B,t-1}, M_{L,t-1}, X_{L,t-1}, \bar{H}_{t-1})$, borrower controls $(c_{B,t}, h_{B,t}, \rho_{B,t}, M^*_{B,t}, H^*_{B,t})$, landlord controls $(c_{L,t}, \rho_{B,t}, M^*_{L,t}, H^*_{L,t})$, saver controls $(c_{S,t}, B_t)$, construction firm controls (L_t, Z_t), and prices $(p_t, q_t, r^*_{B,t}, r^*_{L,t})$ that jointly solve the borrower, landlord, saver, and construction firm problems, as well as the market clearing conditions.

$$
\begin{aligned}
\text{Housing:} \quad & \bar{H}_t = H_{B,t} + H_{L,t} \\
\text{Housing Services:} \quad & \bar{H}_t = h_{B,t} \\
\text{Resources:} \quad & Y_t = c_{B,t} + c_{L,t} + c_{S,t} + Z_t + \delta p_t \bar{H}_t \\
\text{Housing Permits:} \quad & \bar{L} = L_t.
\end{aligned}
$$

The key equation for the model mechanism is the market clearing condition for housing. Combined with the borrower and saver optimality conditions for housing, we obtain inverted demand and supply curves for owner-occupied housing:

$$
\begin{aligned}
p_t^{\text{Demand}}(H_{B,t}) &= \frac{\mathbb{E}_t \left\{ \Lambda_{B,t+1} \left[\bar{\omega}_{B,t} + q_{t+1} + \Big(1 - \delta - (1-\rho_{B,t+1}) \mathcal{C}_{B,t+1}\Big) p_{t+1} \right] \right\}}{1 - \mathcal{C}_{B,t}} \\
p_t^{\text{Supply}}(H_{B,t}) &= \frac{\mathbb{E}_t \left\{ \Lambda_{L,t+1} \left[\bar{\omega}_{L,t} + q_{t+1} + \Big(1 - \delta - (1-\rho_{L,t+1}) \mathcal{C}_{L,t+1}\Big) p_{t+1} \right] \right\}}{1 - \mathcal{C}_{L,t}}.
\end{aligned}
$$

where p_t^{Demand} is the price at which buyers are willing to purchase quantity $H_{B,t}$, and p_t^{Supply} is the price at which sellers are willing to provide quantity $H_{B,t}$ to the market (by market clearing, equivalent to landlords choosing to purchase quantity $\bar{H} - H_{B,t}$ of housing).

These equations map directly into the supply and demand framework of Section 2. The shape of the supply and demand schedules are pinned down by the $\bar{\omega}_{j,t}$ terms. Recall that these terms are defined by $\bar{\omega}_{B,t} = \Gamma_{\omega,B}(1 - H_{B,t}/\bar{H})$ and $\bar{\omega}_{L,t} = \Gamma_{\omega,L}(H_{B,t}/\bar{H})$.

As $H_{B,t}$ increases, so does the fraction of owner-occupied housing. This pushes down $\bar{\omega}_{B,t}$, as the houses obtained become increasingly less suitable for owner-occupation, generating a downward sloping demand curve. At the same time, $\bar{\omega}_{L,t}$ rises as the marginal unit becomes more and more favorable for rental, generating an upward sloping supply curve. At equilibrium, the level of owner-occupied housing $H_{B,t}$ ensures that the two prices p_t^{Demand} and p_t^{Supply} converge, allowing for market clearing. The degree of dispersion in the $\Gamma_{\omega,B}$ and $\Gamma_{\omega,L}$ distributions determine how much the $\bar{\omega}_{j,t}$ terms change with the homeownership rate, governing the slopes of the two curves.

A key variable of interest emerging from equilibrium is the homeownership rate $HOR_{B,t} = \Gamma_{\omega,B}(1 - H_{B,t}/\bar{H})$. Since this variable measures the homeownership among "borrowers" only, for comparison to the data it is necessary to transform this variable slightly. In particular, we compute the overall homeownership rate as

$$HOR_t = \chi_B HOR_{B,t} + (1 - \chi_B)\overline{HOR}$$

where the offsetting term $\overline{HOR}$ is calibrated to ensure that HOR_t is equal to the national homeownership rate in steady state. Since our calibration procedure in Section 7 seeks to match the change in the log homeownership rate in the data, having the correct base rate is important to accurately match this proportional change.

7 Model Quantification

We calibrate our model at quarterly frequency to match several external targets from the literature together with our key empirical moment.

7.1 External Targets and Parameters

Before matching our empirical findings of Section 5, we first describe the more basic portion of our calibration, which maps to more standard moments of the data. The benchmark specification in our preliminary draft turns off landlord credit, so that $M_{L,t} = X_{L,t} = \rho_{L,t} = 0$ for all t. We additionally fix the refinancing rate to be equal to $\rho_t = \bar{\rho} = 0.034$, the steady state refinancing rate from Greenwald (2018). These restrictions are likely to be relaxed in the final version of the paper. The remaining parameters and functional forms are described below and summarized in Table 4.

Table 4: Parameter Values: Baseline Calibration (Quarterly)

Parameter	Name	Value	Internal	Target/Source
		Demographics and Preferences		
Borrower pop. share	χ_B	0.626	N	1998 SCF
Borrower inc. share	s_B	0.605	N	1998 SCF
Landlord pop. share	χ_L	0.020	N	Normalization
Borr. discount factor	β_B	0.970	Y	Borr. VTI (1998 SCF)
Saver discount factor	β_S	0.992	Y	Interest rate = 6.46%
Housing utility weight	ξ	0.200	N	Davis and Ortalo-Magné (2011)
Housing utility elasticity	ε	1.250	N	Piazzesi, Schneider, and Tuzel (2007)
		Ownership Benefit Heterogeneity		
Landlord het. (location)	$\mu_{\omega,L}$	0.000	Y	Normalization
Landlord het. (scale)	$\sigma_{\omega,L}$	0.015	Y	Empirical elasticities (Section 7)
Borr. het. (location)	$\mu_{\omega,B}$	0.008	Y	Avg. Homeownership = 54.82%
Borr. het. (scale)	$\sigma_{\omega,B}$	0.001	Y	Simulated housing boom
		Technology and Government		
New land per period	$\bar{L}$	0.141	Y	Residential inv = 5% of GDP
Land share of construction	φ	0.400	N	Simulated housing boom
Housing depreciation	δ	0.005	N	Standard
Inflation	$\bar{\pi}$	1.008	N	3.22% Annualized
Tax rate	τ	0.200	N	Standard
		Mortgage Contracts		
Refinancing rate	$\bar{\rho}$	0.034	N	Greenwald (2018)
Loan amortization	ν	0.004	N	Greenwald (2018)
LTV Limit	θ^{LTV}	0.850	N	Greenwald (2018)
PTI Limit	θ^{PTI}	0.360	N	Greenwald (2018)
PTI offset (taxes etc.)	α	0.001	N	Greenwald (2018)

Demographics and Preferences To determine the borrower population share, we turn to the 1998 Survey of Consumer Finances. In the model, borrowers are constrained households whose choice of rental vs. ownership is potentially influenced by credit conditions. Correspondingly, we identify a household as a "borrower" if it either (i) owns a home with a mortgage balance greater than 30% of the home's value, or (ii) does not own a home and has less than two months' income in liquid assets. This yields a population share of $\chi_B = 0.626$ and an income share of $s_B = 0.605$. In the absence of good data on the share of landlords, and given that this share is largely an inconsequential normalization in the model, we set the share of landlords to 2%, and assume that landlords do not receive labor income, instead subsisting entirely on their rental earnings.

For preference parameters, we assume a standard consumption weight parameter of $\xi = 0.200$ on housing, following the evidence in Davis and Ortalo-Magné (2011). For the elasticity of the CES consumption aggregator, we follow Piazzesi et al. (2007) and Kaplan et al. (2019) in setting $\varepsilon = 1.250$. We set the borrower discount factor to target the average ratio of home value to income among "borrowers" who own homes in the 1998 SCF, equal to 8.81 (quarterly). We set the saver discount factor to target a nominal interest rate of 6.46%, equal to the average rate on 10-year treasury bills in the immediate pre-boom era (1993 - 1997). To obtain an overall homeownership rate that can be compared to the data, we fix the share of savers who are assumed to own their home to 0.820 which ensures a steady state ownership rate of 65%.

Ownership Cost Heterogeneity. The paper's most novel modeling mechanism relates to heterogeneity in the benefits to borrower and landlord ownership, represented by the distributions $\Gamma_{\omega,B}$ and $\Gamma_{\omega,L}$. We specify each of these as a logistic distribution, so that each c.d.f. is defined by

$$\Gamma_{\omega,j}(\omega) = \left[1 + \exp\left\{-\left(\frac{\omega - \mu_{\omega,j}}{\sigma_{\omega,j}}\right)\right\}\right]^{-1} \qquad j \in \{B, L\}.$$

For the landlord, the location term $\mu_{\omega,L}$ is normalized to zero, while the scale term $\sigma_{\omega,L}$ — the key parameter in the model — is calibrated to match our empirical estimates from Section 5. This calibration procedure is described in detail in Section 7.2. For the borrower distribution, we choose the location parameter $\mu_{\omega,B}$ to attain the correct homeownership rate among "borrowers" in the 1998 SCF (54.82%), and choose the scale parameter $\sigma_{\omega,B}$ to generate a good fit of homeownership rate dynamics in the simulated boom experiment of Section 8.

Technology and Government. For the construction technology, we set the amount of new land permits issued per period to generate residential investment Z_t equal to 5% of total output in the steady state. For the land weight in the construction func-

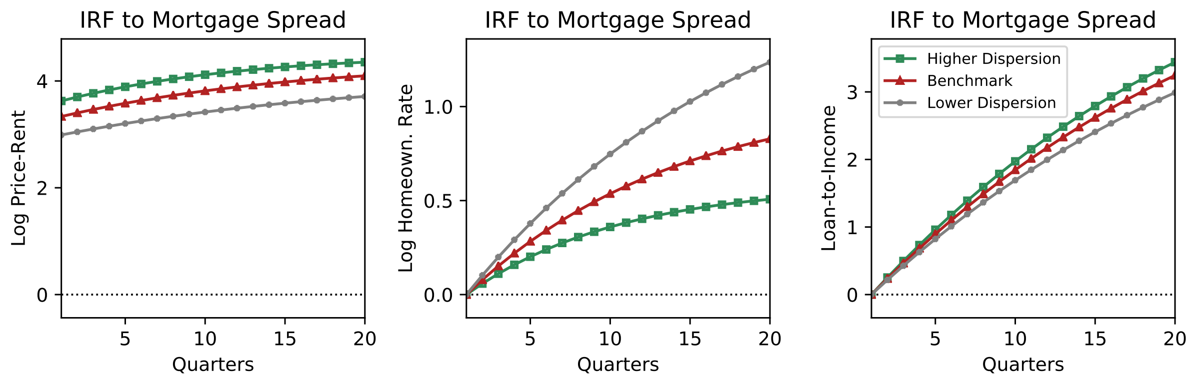

Figure 4: Model Response to our Identified Credit (CLL) Shock by $\sigma_{\omega,L}$

tion φ, we choose 0.400 to attain a good fit of the level of residential investment in our simulated housing boom experiments. We set housing depreciation and the tax rate to standard values, and set inflation to be equal to the average 10-year inflation expectation in the pre-boom era (1993-1997) following Greenwald (2018).

Mortgage Contracts. For the mortgage contract parameters, we follow Greenwald (2018), who provides a detailed calibration for this mortgage structure.

7.2 Calibration of Landlord Heterogeneity to Our Empirical Results

The key parameter governing the model's quantitative behavior is the dispersion in the cost of landlord ownership, which determines the slope of the housing supply curve. We calibrate this parameter so that the model is able to reproduce our empirical findings on the relative sensitivity of the price-to-rent ratio and homeownership rate to a credit supply shock. Since our instrument exploits a change in the conforming loan limit, which changes the cost at which a treated household can obtain credit, we choose as our model equivalent an impulse response with respect to the mortgage spread, $\Delta_{B,t}$. As a change in the conforming loan limit is expected to persist indefinitely, we assume a near permanent shock with quarterly persistence 0.995.

Figure 4 shows this impulse response for three possible calibrations of our dispersion parameter, $\sigma_{\omega,L}$. Our benchmark calibration of $\sigma_{\omega,L} = 0.015$ is chosen to match our target of a five-to-one ratio between the elasticities of the price-to-rent ratio and the homeownership rate, with 20Q growth in these variables of 4.10 and 0.83 , respectively. The other two series show paths for alternative calibration choices: the "Higher

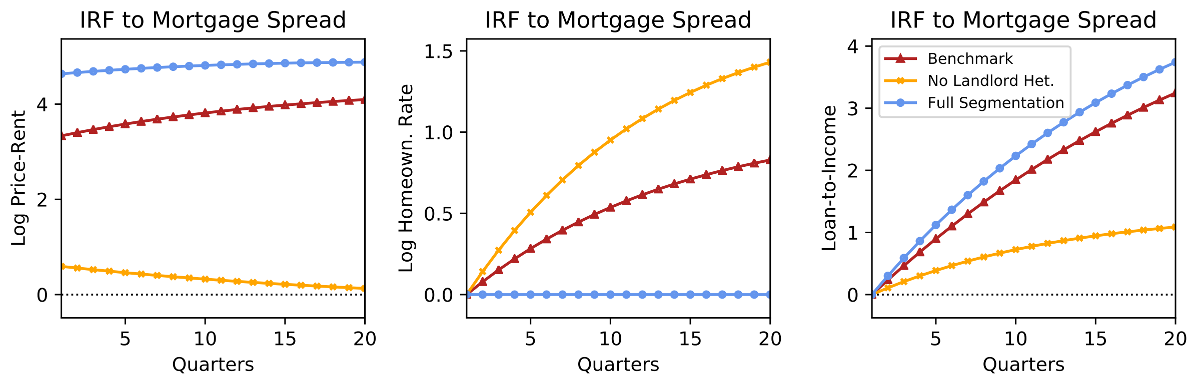

Figure 5: Model Response to our Identified Credit Shock, Benchmark vs. Polar Cases

Dispersion" series doubles the dispersion to $\sigma_{\omega,L} = 0.030$, while the "Lower Dispersion" series halves the dispersion to $\sigma_{\omega,L} = 0.0075$. The figure shows that these alternative calibrations lead to substantially different ratios for these elasticities, largely due to the high sensitivity of the homeownership rate elasticity to this parameter. In particular, the "Higher Dispersion" series generates a 20Q ratio of 8.6 (4.35 vs. 0.51 , respectively) while the "Lower Dispersion" series generates a 20Q ratio of 3.00 (3.71 vs. 1.24 , respectively). Overall, this comparison illustrates that $\sigma_{\omega,L}$ can be identified using this procedure.

As a final note, we point out that this calibration procedure should be valid even in a richer model in which landlords use credit. The key point is that our identified shock is specifically to borrowing conditions for *households*, through the conforming loan limit. As a result, this procedure should pin down the slope of the supply curve, even if at other times the supply curve itself shifts due to expansions or contractions in landlord credit.

7.3 The Role of Credit in Our Calibrated Model

Before we examine the role of credit in the boom and bust of the 2000s, it is instructive to consider to consider where our calibrated response to this credit shock compares to the extremes explored in the literature: full segmentation ($\sigma_{\omega,L} \to \infty$) and frictionless rental markets ($\sigma_{\omega,L} \to 0$).

To this end, Figure 5 compares the same impulse response (a near-permanent shock to mortgage spreads, $\Delta_{B,t}$) For comparison, the path labeled "No Landlord Het." displays the same impulse response in an alternative economy with no landlord heterogeneity ($\sigma_{\omega,L} \to 0$). Contrasting these two series shows that our benchmark calibration implies a substantial departure, generating large positive increase in price-to-rent ratios compared to the no heterogeneity world, but displaying just over half of the increase in the homeownership rate (0.83% vs. 1.43%), in line with the intuition of Section 2.

Importantly, the benchmark economy also generates much higher credit growth, as shown in the third panel of Figure 5. This is a direct consequence of the increased price response due to the steeper housing supply curve under landlord heterogeneity. Although the "No Landlord Het." economy sees more renters become mortgage-holding owners following this shock, generating a larger increase in credit at the extensive margin, this effect is overwhelmed by the increase in house value in the benchmark economy, whose more valuable collateral boosts the size of new loans. Overall, the ratio of mortgage debt to total borrower income grows by roughly three times more in the benchmark vs. alternative economy (3.24% vs. 1.09%), implying that models without heterogeneity may seriously understate the impact of credit shocks on household leverage.

In contrast, the benchmark model is much closer to the other extreme case, labeled "Full Segmentation," which has a fixed homeownership rate ($\sigma_{\omega,L} \to \infty$). In fact, the benchmark model generates 84% and 87% of the 20Q rise in price-to-rent and loan-to-income, respectively, found under full segmentation. These results demonstrate that matching our empirical findings requires substantial frictions, with house price and credit dynamics that are more similar to full segmentation than to a frictionless rental market. Of course, the benchmark model does deliver meaningful movements in the homeownership rate that are precluded in the "Full Segmentation" model, implying that a departure from this extreme is still essential to match the full set of dynamics.

7.4 Saver Housing Demand

A key assumption in our modeling framework is that savers do not consume housing, eliminating any possible transactions between savers and borrowers. An alternative model where both groups consume housing and can groups to frictionlessly trade houses would exhibit dramatically dampened responses of house prices to credit. The reasoning is very similar to that for frictionless rental markets: the willingness of unconstrained savers, who do not rely on credit, to buy and sell housing at the margin, stabilizes the price as credit shifts.

While our modeling assumption is extreme, housing markets are much more segmented in reality than in our model. In our model, as with nearly all macro-housing models, the housing good is perfectly divisible in our model, potentially allowing savers to buy up units of housing and agglomerate them into their existing home. The willingness of unconstrained savers to frictionlessly add housing services to their consumption bundle therefore provides a relatively elastic margin of adjustment. This margin is likely much more restricted in reality, as housing in the real world occupies vast spectra of quality levels and geographic locations, implying that households cannot easily enjoy more service flows by purchasing additional housing. Specifically, unconstrained agents likely gain little service flow from second homes in struggling areas, making them much less willing to purchase housing at the extensive margin compared to the model world in which additional houses could essentially be absorbed into the primary residence.

Importantly, our calibration procedure ensures that our results should not be severely biased by this modeling assumption. The reason is that our empirical estimates measure the slope of the *overall* supply of housing to borrowers, which is a combination of the supply from landlords and unconstrained savers. If unconstrained savers are transacting with borrowers in the data, this should dampen the house price response to credit, while leaving the homeownership rate response unchanged. Although our model ignores this mechanism, incorporating it would require us to calibrate even more extreme rental market frictions to reproduce our empirical findings, leading to similar overall results.

8 Model Results

Now that the model has been calibrated to match our empirical results, we can ask the core quantitative question of our paper: what role did credit play in the housing boom? To do this, we simulate a realistic relaxation of credit standards and evaluate the model's implications for the evolution of debt and house prices. For simulated credit relaxation, our baseline experiment follows KMV in relaxing LTV limits from

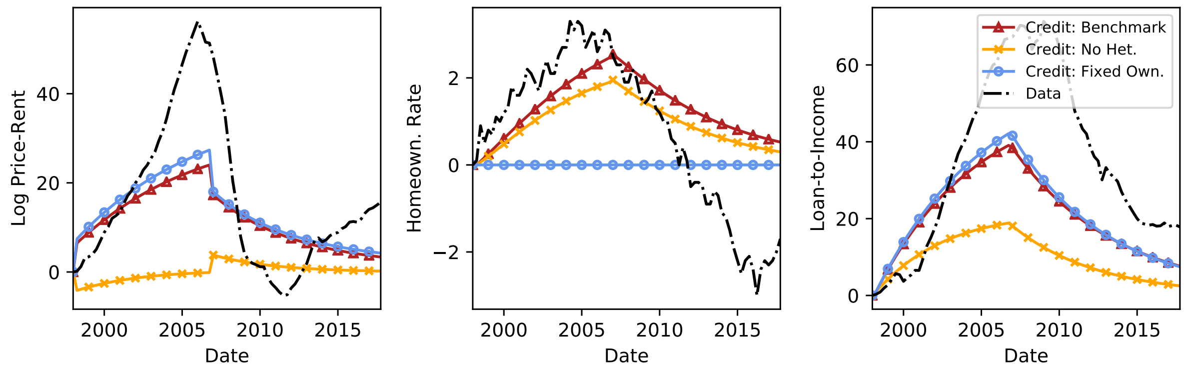

Figure 6: Credit Relaxation Experiment

85% to 99%, and relaxing PTI limits from 36% to 65%. We impose this relaxation as an unexpected but permanent change in 1998 Q1. The new standards are left in place until 2007 Q1, at which time they unexpectedly and permanently revert to their original values. The model responses are then computed as nonlinear perfect foresight paths.

The results of this experiment are shown in Figure 6. To highlight the role of landlord heterogeneity, we again plot the responses in our Benchmark model against those of an alternative model with no landlord heterogeneity. The results show that our calibrated level of heterogeneity is large enough to deliver a large price response in the Benchmark model, accounting for 47% of the peak rise in price-to-rent ratios observed in the boom. This stands in sharp contrast to the model without landlord heterogeneity, where the same credit relaxation explains 0% of the peak growth in price-to-rent ratios, as landlords are able to completely satisfy the increase in demand, preventing a rise in prices. As before, the house price dynamcis Benchmark model are much closer to the "Full Segmentation" model, where this credit relaxation would account for 53% of the observed rise in price-to-rent ratios. Overall, these results indicate that a realistically calibrated rental market still delivers an economically important response of house prices to a relaxation of credit conditions similar to that observed in the 2000s boom.

This finding for house prices also has important implications for credit growth. While credit standards are loosened equally along both paths of Figure 6, credit growth over the boom is more than twice as high in the Benchmark economy as in the alternative economy, explaining 58% vs. 28% of the observed rise. This additional credit growth is a direct consequence of the larger house price appreciation in the Benchmark

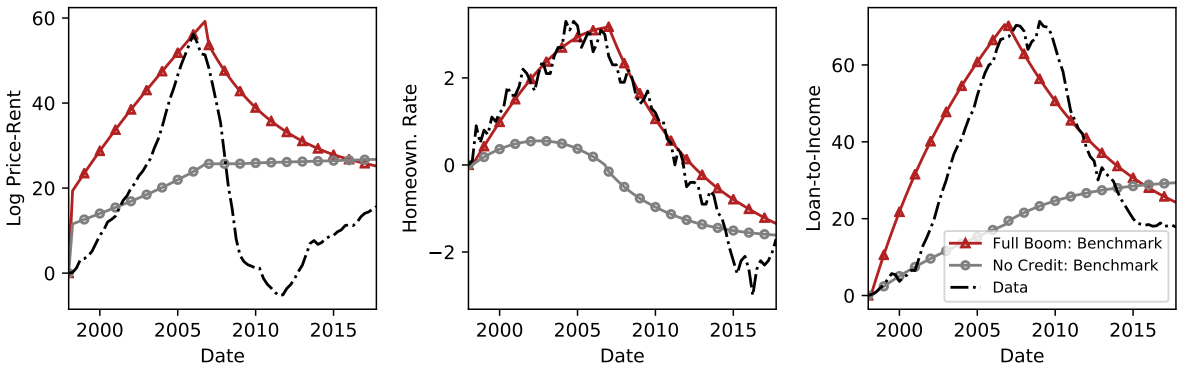

Figure 7: Complete Boom Experiment

economy, which increases the value of housing collateral, and allows larger loans for a given maximum LTV ratio. As a result, the same credit loosening leads to much more levered households in the Benchmark economy when credit conditions return to baseline. Again, the Benchmark path very close to the Full Segmentation path, which also exhibits a large house price rise, and therefore delivers credit growth corresponding to 63% of the observed rise.

Although our Benchmark model indicates that a credit expansion played an important role in driving the boom, it clearly leaves room for other factors to play important roles. To explore what our results indicate in a more comprehensive simulation of the boom, we first incorporate a 2ppt fall in mortgage spreads, assumed to be permanent. Like our credit expansion, this causes an outward shift of housing demand, which, given our estimated rental frictions generates a large additional increase in house prices. Specifically, combining our credit relaxation and the fall in rates (not shown) can explain 85% of the rise in price-to-rent ratios and 90% of the rise in loan-to-income ratios. Our results therefore indicate that realistically calibrated rental frictions, combined with a relaxation of credit standards and a decline in credit costs can explain the vast majority of the housing and credit cycle.

To complete our explanation of the boom, we add overoptimistic house price expectations to generate a rise in price-to-rent ratios equal to that observed in the data.[15] Unlike the credit experiments, this shift in expectations shifts the supply curve instead of moving along it. We find that matching the overall rise in house prices requires that

[15]Overoptimistic house price expectations are modeled as expected increases in the housing utility parameter ξ. Agents believe this parameter will increase in 2007Q1, but at that time it is instead revealed to stay at its original level.

agents believe in a future increase in ξ, the expenditure share on housing services, of 40%. Given the size of this required expected growth, we note that part if not all of this supply shift required to fit the data could alternatively be generated by shifts in landlord access to credit, to be explored in future work.[16]

To match the bust, we impose a further 3% fall in interest rates, which is also reflected in a 3% fall in the landlord discount rate, consistent with the entry of yield-seeking financial firms into the single family rental market, as well as a 10% decline in LTV and PTI limits, consistent with tightening credit standards.[17] Overall, these assumptions generate a reasonably good fit of the dynamics of the boom and bust, with two main exceptions: (i) house prices adjust much less sluggishly in our model than in the data, as is typically found in models lacking a rich set of search frictions as in e.g., Guren (2018); and (ii) our model "bust" is much more gradual in the model relative to the data, as we lack the foreclosures and financial market features that transformed the housing crash into a global crisis.

To isolate the role of credit conditions in this simulated boom-bust, we then remove the simulated credit expansion, while leaving all the other factors in place, to generate the series labeled "No Credit: Benchmark." These results indicate that removing the credit expansion from the rest of the boom would have reduced the overall rise in price-to-rent ratios by 57% and in loan-to-income ratios by 74%. These shares are larger than the shares explained by relaxing credit in isolation (Figure 6), because in the Benchmark model, loose credit also amplifies the other components of the boom, particularly so for overoptimistic house price expectations.

To summarize our results, our calibrated model implies an important role for credit conditions in explaining the housing and credit cycles observed in the 2000s boom-bust, with a relaxation of credit standards explaining roughly half the rise in price-to-rent ratios, with house price dynamics that are closer to the extreme of full segmentation than to a frictionless model with no landlord heterogeneity.

[16] A shift in landlord risk premia, similar to the results in Favilukis et al. (2017), could also contribute to this shift.

[17] This is best interpreted as increasing standards for credit scores preventing a fraction of the population from obtaining credit at the extensive margin, rather than a decline in maximum LTV and PTI ratios at the intensive margin.

9 Conclusion

More than a decade after the Great Recession there is still a lack of consensus about the role of credit supply in explaining house prices and price-to-rent ratios in the boom and bust. In this paper, we argue that this is because the literature has focused on two polar cases. In the first, landlords are unable to convert homes from owner-occupied to renter-occupied and savers cannot step in to purchase (sell) more housing when constrained households cut (raise) their housing demand. This means that all changes in housing demand show up in house prices and homeownership rates are stable. At the other extreme, models feature either deep-pocketed landlords who can frictionlessly convert between owner-occupied and rental housing, or unconstrained savers with elastic demand for housing, who completely absorb changes in borrower housing demand due to credit conditions. These features imply that changes in credit supply lead to large changes in homeownership rates but no movement in price-to-rent ratios.

We generalize these polar cases to examine intermediate levels of rental frictions, which we view as realistic. We argue that a key sufficient statistic for determining where reality falls on the spectrum between these two extremes is the relative causal effect of credit on the price-to-rent ratio vs. the homeownership rate. We show in a new data set using instrumental variables methods that credit supply shocks cause a significant increase in price-to-rent ratios and a more muted and statistically insignificant homeownership response. We propose five-to-one as a conservative estimate of the ratio of the two casual effects. When we calibrate a model to match this ratio, we find that credit supply can explain roughly half of the boom and bust in house prices and price-to-rent ratios. Relative to our polar cases, the calibrated model displays house price dynamics that are close to those under perfect segmentation, implying large frictions in rental markets.

Our work highlights the importance of assumptions about rental markets and the elasticity of saver demand for macro models of the housing market. These model features are often overlooked but are critical for many important results. We hope that our findings motivate future work to use and develop intermediate models in place of either polar assumption. We also highlight the use of identified credit supply shocks and a novel empirical moment for calibrating macroeconomic models of the housing market. We hope that future work will improve on our estimates of the relative causal effect of credit supply on price-to-rent ratios and homeownership rates, and use these identified moments to improve the calibration of macro-housing models.

References

CHEN, V. W. (2018): "Fiscal Multipliers and Regional Reallocation," Unpublished Paper. Boston University.

DAVIS, M. A. AND F. ORTALO-MAGNÉ (2011): "Household Expenditures, Wages, Rents," *Review of Economic Dynamics*, 14, 248–261.

FAVARA, G. AND J. IMBS (2015): "Credit Supply and the Price of Housing," *American Economic Review*, 105, 958–992.

FAVILUKIS, J., S. C. LUDVIGSON, AND S. VAN NIEUWERBURGH (2017): "The Macroeconomic Effects of Housing Wealth, Housing Finance, and Limited Risk Sharing in General Equilibrium," *Journal of Political Economy*, 125, 140–223.

GARRIGA, C. AND A. HEDLUND (2017): "Mortgage Debt, Consumption, and Illiquid Housing Markets in the Great Recession," Working Paper, Federal Reserve Bank of St. Louis and University of Missouri.

——— (2018): "Housing Finance, Boom-Bust Episodes, and Macroeconomic Fragility," Working Paper, Federal Reserve Bank of St. Louis and University of Missouri.

GARRIGA, C., R. MANUELLI, AND A. PERALTA-ALVA (2019): "A Macroeconomic Model of Price Swings in the Housing Market," *American Economic Review*, Forthcoming.

GETE, P. AND M. REHER (2018): "Mortgage Supply and Housing Rents," *Review of Financial Studies*, 31, 4884–4911.

GREENWALD, D. (2018): "The Mortgage Credit Channel of Macroeconomic Transmission," MIT Sloan Research Paper No. 5184-16.

GUREN, A. AND T. MCQUADE (2018): "How do foreclosures exacerbate housing downturns?" Working Paper, Boston University and Stanford University.

GUREN, A. M. (2018): "House Price Momentum and Strategic Complementarity," *Journal of Political Economy*, 126, 1172–1218.

GUREN, A. M., A. KRISHNAMURTHY, AND T. J. MCQUADE (2018): "Mortgage Design in an Equilibrium Model of the Housing Market," Working Paper, Boston University and Stanford University.

JORDÀ, Ò. (2005): "Estimation and Inference of Impulse Responses by Local Projections," *American Economic Review*, 161–182.

JUSTINIANO, A., G. E. PRIMICERI, AND A. TAMBALOTTI (2015a): "Credit Supply and the Housing Boom," Tech. rep., National Bureau of Economic Research.

——— (2015b): "Household Leveraging and Deleveraging," *Review of Economic Dynamics*, 18, 3–20.

KAPLAN, G., K. MITMAN, AND G. L. VIOLANTE (2019): "The Housing Boom and Bust: Model Meets Evidence," .

LOUTSKINA, E. AND P. E. STRAHAN (2015): "Financial Integration, Housing, and Economic Volatility," *Journal of Financial Economics*, 115, 25–41.

PIAZZESI, M., M. SCHNEIDER, AND S. TUZEL (2007): "Housing, Consumption and Asset Pricing," *Journal of Financial Economics*, 83, 531–569.

RAMEY, V. A. (2016): "Macroeconomic Shocks and Their Propagation," *Handbook of Macroeconomics*, 2, 71–162.

RAMEY, V. A. AND S. ZUBAIRY (2018): "Government Spending Multipliers in Good Times and in Bad: Evidence from U.S. Historical Data," *Journal of Political Economy*, 126, 850–901.

SAIZ, A. (2010): "The Geographic Determinants of Housing Supply," *The Quarterly Journal of Economics*, 125, 1253–1296.

STOCK, J. H. AND M. W. WATSON (2018): "Identification and Estimation of Dynamic Causal Effects in Macroeconomics Using External Instruments," *The Economic Journal*, 128, 917–948.