Do Credit Conditions Move House Prices?

Daniel L. Greenwald* and Adam Guren†‡

September 30, 2019

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 most of the house price cycle. We show that the key difference behind these disparate results is the extent to which credit insensitive agents such as landlords and unconstrained savers absorb credit-driven demand, which depends on the degree of segmentation in housing markets. We develop a model with frictional rental markets which allows us to consider cases in between the extremes of no segmentation and perfect segmentation typically assumed in the literature. 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 the degree of segmentation. We estimate this moment using regional variation in credit supply and use it to calibrate our model. Our results reveal that rental markets are highly frictional and close to fully segmented, which implies large effects of credit on house prices. In particular, changing credit conditions can explain between 28% and 47% of the rise in price-rent ratios over the boom.

*MIT Sloan School of Management, dlg@mit.edu

†Boston University and NBER, guren@bu.edu

‡We would like to thank Joe Vavra, Monika Piazzesi, Tim McQuade, Jaromir Nosal, and seminar participants at the 2019 AEA Annual Meeting, Duke University, SED, and SITE for helpful comments. We would like to thank CBRE Economic Advisers (particularly Nathan Adkins and Alison Grimaldi) 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 and the effectiveness of macroprudential policy tools, 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 the majority of the movements in house prices in the 2000s.[1] By contrast, other papers, such as 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 these models 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 the degree to which credit insensitive agents such as landlords and unconstrained savers absorb credit-driven demand by constrained agents. This in turn depends on the degree of segmentation between credit insensitive and credit sensitive agents in housing markets. Many models assume either perfect segmentation or perfect integration for tractability, and this assumption plays a crucial role in the findings.

The importance of segmentation is clearest between rental and owner-occupied housing. In models with complete integration, 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. Credit-induced shifts in demand for homes by constrained agents move the homeownership rate but have 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 con-

[1] Favilukis et al. (2017) find that 60% of the rise in price to rent ratios can be explained by credit alone and all of the rise can be explained by a combination of credit and business cycle shocks. Landvoigt, Piazzesi, and Schneider (2015), Greenwald (2018), Guren, Krishnamurthy, and McQuade (2018), Garriga, Manuelli, and Peralta-Alva (2019), Garriga and Hedlund (2017), Garriga and Hedlund (2018), Justiniano, Primiceri, and Tambalotti (2015a), and Liu, Wang, and Zha (2019) also analyze models in which credit conditions play a key role in the boom and bust. Conversely, Kiyotaki, Michaelides, and Nikolov (2011) study a model in which credit conditions play a limited role.

trast, in models with segmentation — because homes differ in their suitability for renting — landlords cannot step in to buy houses when credit conditions change. Shifts in credit supply then 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.

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 degree of segmentation in housing markets and thus the response of house prices to credit. T his is because this ratio corresponds to the slope of the 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. Importantly, the slope of the supply curve around the economy's equilibrium does not depend on whether the marginal buyer is a landlord or unconstrained saver. This slope consequently 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. To this end, the first part of our paper provides new empirical estimates of the causal effect of credit supply on prices, rents, and homeownership rates. 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 two off-the-shelf instruments. The first and primary empirical approach 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). The second empirical approach uses differential city-level exposure to a sudden expansion in the private label securitization market following Mian and Sufi (2019). Both sets of estimates indicate that shocks to credit supply increase prices, rents, and the price-to-rent ratio, but have either no statistically significant effect or a much smaller effect on the homeownership rate. In terms of point estimates, both approaches find that price-to-rent ratios respond roughly 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-occupied housing, as well as a continuum of intermediate cases. 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 between 28% and 47% 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 in housing markets. Our benchmark calibration generates house price dynamics that are 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.

Our baseline model makes two stark assumptions: That landlords do not use credit and that the saver housing stock is entirely segmented from the borrower housing stock.

We relax each of these assumptions in turn. When landlords require credit, a credit supply expansion that also affects landlords shifts supply out, leading to a larger price-to-rent ratio response and a smaller homeownership rate response than under our baseline results. Intuitively, when landlords have more access to credit, they can more easily build rental housing, rents fall, and more households are induced to rent. Thus if anything our results without landlord credit are a lower bound on the price response to a credit supply shock.

A similar logic to our main result about segmentation holds for the degree of segmentation between unconstrained savers and constrained households. If these households compete for a single homogenous housing good, then the marginal valuation of a unit by unconstrained savers pins down the price, generating an elastic supply curve and limited response of price to credit. Intuitively, the savers act as the landlords in the rental markets case. By contrast, if credit constrained agents operate mainly in lower-quality segments of the housing stock (Landvoigt et al. (2015)) or tend to buy smaller "starter homes" (Ortalo-Magné and Rady (2006)) and housing is non-divisible, then the housing demand of unconstrained unconstrained savers may be quite segmented from that of the credit constrained agents. In such a model, credit again plays an important role for the segments of the market in which constrained agents are the marginal buyer (Landvoigt et al. (2015)). As with rental markets, the degree of housing market segmentation between unconstrained and constrained agents is crucial.

Related Literature. Our paper relates to several literatures. Empirically, our analysis builds on prior analyses of the causal effect of credit and interest rates on house prices including Glaeser, Gottlieb, and Gyourko (2012), Adelino, Schoar, and Severino (2012), Favara and Imbs (2015), Loutskina and Strahan (2015), Di Maggio and Kermani (2017), Mian and Sufi (2019), and Gete and Reher (2018). These papers, however, cannot answer what fraction of the boom and bust can be explained by credit unless the quasi-random variation they use corresponds exactly to the shocks that drove the boom and bust. We build on this literature by showing that the ratio of the causal effect of credit on price-to-rent to the causal effect on the homeownership rate for any credit variation can be used as a moment to identify structural elasticities that can be used to calculate the effect of credit on house prices for *any* set of shocks, including those that correspond to the 2000s boom and bust.

In terms of applied theory, our paper relates to papers that study the effect of credit supply on house prices such as Favilukis et al. (2017), Kaplan et al. (2019), Kiyotaki et al. (2011), Greenwald (2018), Guren et al. (2018), Garriga et al. (2019), Garriga and Hedlund (2017), Garriga and Hedlund (2018), Justiniano et al. (2015a), Liu et al. (2019), and Huo and Rios-Rull (2016). The most closely related paper is Landvoigt et al. (2015), who use an assignment model calibrated to micro data to study endogenous segmentation between constrained and unconstrained homeowners who sort into homes of different quality and find that credit is important in explaining the boom at the bottom of the quality distribution. Our model of borrower-saver segmentation is more reduced form, but our tractable approach allows us to add segmentation due to limited rental conversion and embed the housing market in a more complete general equilibrium model that provides for a richer set of counterfactuals. Nonetheless, we see our results as highly complementary to Landvoigt et al.'s.

Our paper also relates to work on macroprudential policy. Because mortgage credit dominates household balance sheets, many macroprudential policies only work if credit affects house prices. Similarly, *ex-post* debt reduction and foreclosure policies (Guren and McQuade (2018), Mitman (2016), Agarwal, Amromin, Ben-David, Chomsisengphet, Piskorski, and Seru (2017a), Agarwal, Amromin, Chomsisengphet, Landvoigt, Piskorski, Seru, and Yao (2017b), Hedlund (2016)) and mortgage design (Guren et al. (2018), Greenwald, Landvoigt, and Van Nieuwerburgh (2017), Campbell, Clara, and Cocco (2018), Piskorski and Tchistyi (2017)) have additional bite if they affect house prices.

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 extends the

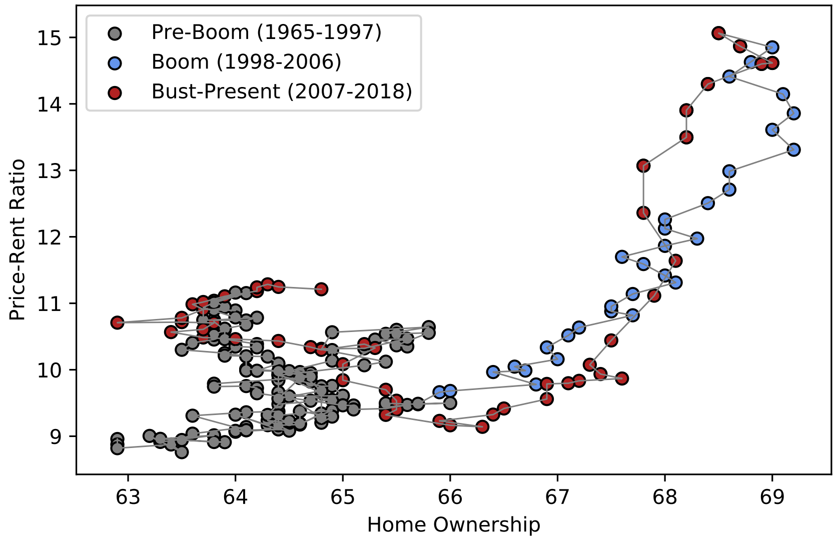

Figure 1: Price-Rent Ratio vs. Homeownership Rate

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).

model to include saver housing demand and credit for landlords. Section 10 concludes.

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,

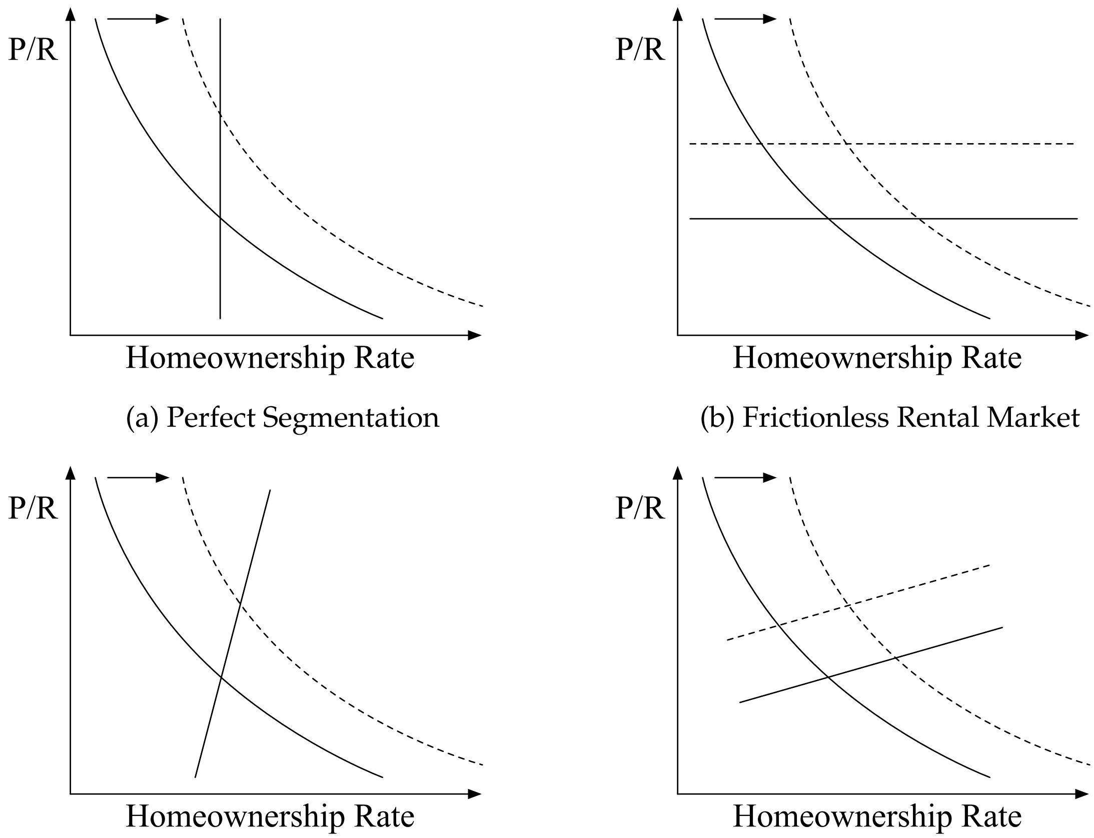

(a) Perfect Segmentation

(b) Frictionless Rental Market

(c) Intermediate Case 1

(d) Intermediate Case 2

Figure 2: Supply and Demand

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 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 (horizontal) 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.[2] 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 multi-

[2]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.

family 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 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.[3]

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.[4] 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 repeat 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

[3]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.

[4]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 as the average change in rents for identical units in the same buildings. CBRE does not use the standard repeat sales methodology because rent data is available for most buildings continuously, so accounting for many periods of missing prices is unnecessary.

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.

For our second instrument, we use an alternate data set for a broader set of 246 CBSAs. House prices come from the FHFA for this broader set of CBSAs, and homeownership rates come from the American Community Survey (ACS) 1-year estimates from the Census. The ACS data is only available from 2005 to 2017, so this is a shorter panel.

4 Empirical Approach

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.

We use two different off-the-shelf identification approaches. While each has its flaws, they provide similar results. We also plan to incorporate additional instruments in a future draft.

4.1 Main Instrument: Loutskina and Strahan

Our main 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). Adelino et al. (2012) also use a similar empirical strategy. 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.[5] 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.[6]

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.

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

[6]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 particular, one might be concerned that cities with higher prices tend to be more exposed to national business cycles. In the Appendix, we conduct a number of robustness checks to assuage some concerns, including adding city-level income shocks and time-varying controls for the industrial structure. Nonetheless, we cannot control for all of the potential issues with our primary instrument, which is why we pursue alternate instruments.

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.[7] 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).

[7] 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.

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 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 Alternate Approach: Mian and Sufi

Our second approach is to use differential city-level exposure to the 2003 expansion in private label securitization (PLS) to identify the effect of credit supply on prices and homeownership rates based on Mian and Sufi (2019). Mian and Sufi build on evidence from Justiniano, Primiceri, and Tambalotti (2017) of a sudden, sizable, and persistent expansion in PLS markets in late 2003 that persists until the crash. Mian and Sufi argue that the PLS expansion had a larger effect on lending by mortgage lenders that rely on non-core deposits to finance mortgages, so they define a bank-level measure of the ratio of non-core liabilities to total liabilities (NCL). High NCL banks, which are funded less by deposits, should have a greater exposure to the PLS expansion. They show that high NCL banks did in fact expand their lending more after following roughly parallel trends prior to 2002. Mian and Sufi then create a measure of CBSA-level exposure to high NCL lenders, which is equal to the origination-weighted average of lender-level NCLs in a CBSA based on 2002 originations. Mian and Sufi argue that the city-level NCL exposure satisfies an exclusion restriction and is a valid instrument for credit.

We use the Mian and Sufi NCL instrument as an alternate empirical approach. While NCL exposure is not a perfect instrument, it provides an alternate exclusion restriction that can be used to corroborate our main results. It also isolates a somewhat different local average treatment effect since it is a non-prime credit supply shock, while the Loutskina-Strahan instrument is a prime credit supply shock.

Because the PLS market expanded once in 2002, we cannot use the panel data techniques we used for our main instrument. Consequently, we adopt Mian and Sufi's specification and examine the reduced form – that is the effect of the NCL share interacted with years on outcomes – rather than using the NCL share as an instrument. In particular, we regress:

$$\log(outcome_{i,t}) = \xi_i + \psi_t + \sum_{t \neq base} \beta_k NCLShare_{i,2002} + \epsilon_{i,t}, \tag{5}$$

where *base* is an excluded base year and the entire regression is weighted by housing units in 2000.

This identification strategy requires more CBSAs than the 57 for which we have Census HVS homeownership rates and Torto-Wheaton rents, so we use ACS data for homeownership rates instead. This gives us 246 CBSAs, but it means that we can only use prices as outcomes.[8] Additionally, we can only start the homeownership rate regressions in 2005 when the ACS begins. We thus cannot use the base year of 2002 that Mian and Sufi use. We use 2013 as the base year as the impulse response for prices returns to its 2002 level in 2013, and our results are robust to using peak-to-trough changes rather than a base year.

[8] We do not use the ACS homeownership rates for the main instrument because the ACS begins in 2005 and most of the variation in the conforming loan limit comes before 2005. The ACS does have rents, but they are average rents and very sticky so a price-to-rent ratio constructed with ACS data looks like the same regression using prices as an outcome

4.5 Future Work: Additional Instruments

The variation in our instruments 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 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 Loutskina-Strahan 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.

[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 unless we use the ACS, which starts in 2005.

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.

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 insignificant.[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

[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.

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.

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 tenants 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.

Our results for the effect of credit on house prices are similar to those found in the literature, such as Glaeser et al. (2012), Adelino et al. (2012), Favara and Imbs (2015), and Di Maggio and Kermani (2017). Our results are also related to the findings of these papers on the relationship between interest rates and house prices.

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 large decrease in the price-to-rent ratio. If one compares their coefficients on the price-to-rent ratio to the homeownership rate, one obtains a ratio of 85 to one, while we calibrate to a five-to-one ratio. Our results may differ for three reasons. First, Gete and Reher use denial rates as a measure of credit supply instead of 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.[12]

5.2 Panel Local Projection IV Impulse Responses

Figure 3 shows the impulse responses using panel local projection instrumental variables.[13] 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.[14] 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] In a robustness check, Gete and Reher use the Loutskina-Strahan instrument prior to the bust and find no significant effect of credit on rents. This is similar to what we find for measures of average rents rather than new rents.

[13] 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.

[14] 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

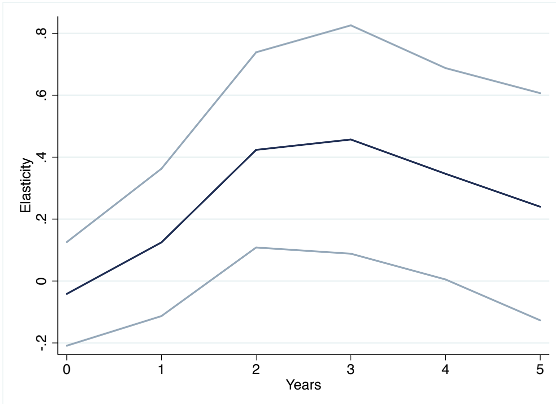

(a) Price/Rent

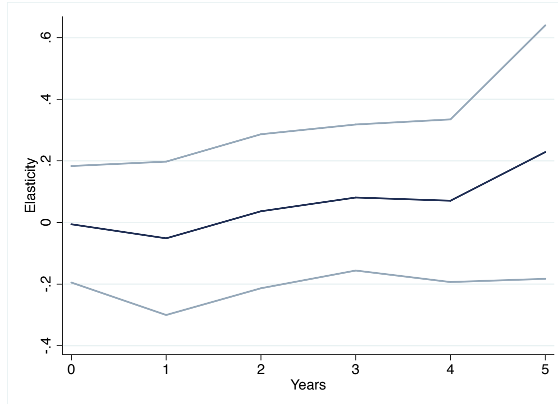

(b) Homeownership Rate

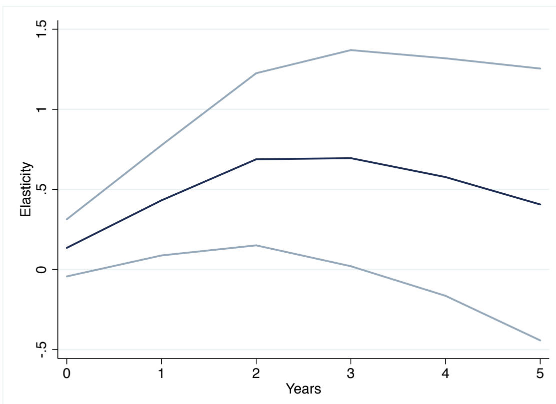

(c) House Price

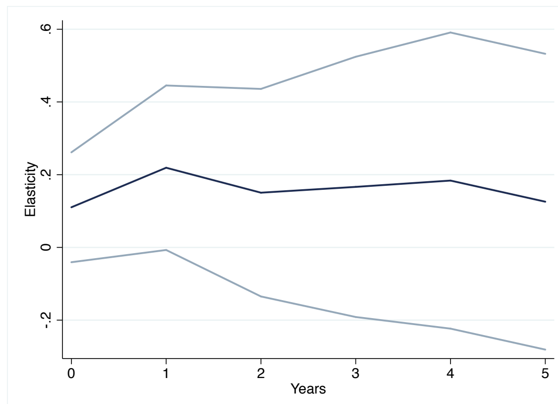

(d) Rent

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.

Figure 4: Mian Sufi NCL Instrument

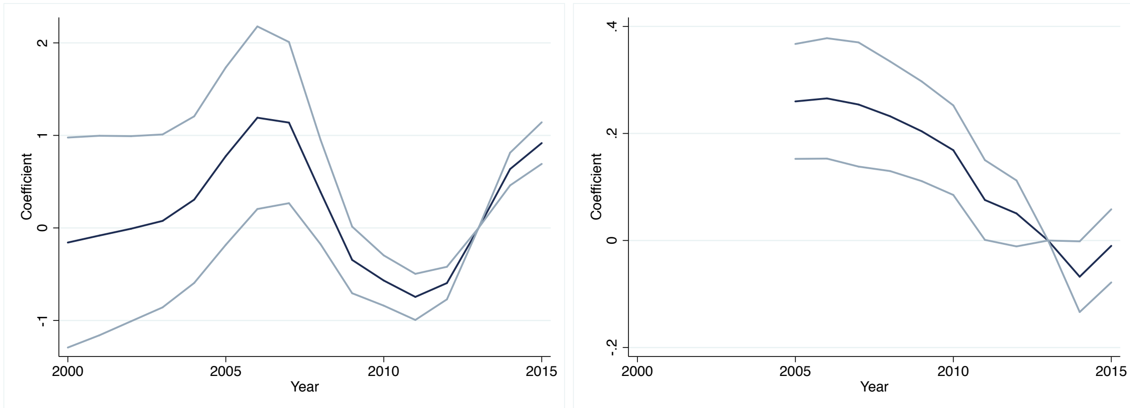

(a) House Prices (b) Homeownership Rate

Notes: 95 % confidence interval shown by gray bands. The figure shows estimates of the effect of a city's NCL share on each outcome based on estimating equation 5. All standard errors are robust.

5.3 Alternate Instrument

Figure 4 shows the year β_k coefficients for the Mian and Sufi NCL share based on equation 5. As a reminder, the PLS market expansion is in 2002, and we have normalized 2013 to be the base year. Panel A shows a zero effect on prices prior to 2002 and then a hump-shaped impulse response in line with our previous results. Panel B shows a positive and statistically-significant effect of the NCL share on homeownership rates beginning in 2005 that also mean reverts over time. Examining the coefficients reveals that the effect of house prices is 4.5 times larger than the response to homeownership. If we were to abandon the 2013 base year and instead use peak to trough changes, we would get a 5.8 times larger response to homeownership rates.

Both of our empirical strategies may have flaws. However, the fact that two separate identification strategies that provide credit supply variation in different segments of the mortgage market provide us with a roughly five-to-one ratio for the response of price-to-rent ratios and homeownership rates provides us with confidence on this key moment. We hope that the literature continues to make progress on precisely estimating this key relative causal effect in the future.

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 and savers both have log preferences over a Cobb-Douglas aggregator of nondurable consumption and housing services:

$$U_B = \sum_{t=0}^{\infty} \beta_B^t \log\left(\hat{c}_{B,t}^{1-\xi} \hat{h}_{B,t}^{\xi}\right)$$
$$U_S = \sum_{t=0}^{\infty} \beta_B^t \log\left(\hat{c}_{S,t}^{1-\xi} \hat{h}_{S,t}^{\xi}\right)$$

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 $j \in \{B, S\}$. Importantly, however, savers are restricted to always demand the fixed quantity of housing $\bar{H}_S$. This is equivalent to assuming a completely segmented housing market, in which savers and borrowers consume different types of housing (e.g., live in different neighborhoods, occupy different quality tiers). 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

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

which correspond to the interpretation of landlords as a foreign-owned, profit-maximizing firm, as in e.g., Kaplan et al. (2019).

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:

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

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.[15]

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

[15]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.

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 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:

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

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

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

while the landlord's laws of motion are

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

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

$$c_{S,t} \leq \underbrace{(1-\tau) y_{S,t}}_{\text{after-tax income}} - \underbrace{(B_t - R_{t-1} B_{t-1})}_{\text{net bond purchases}} - \underbrace{p_t \left(H^*_{S,t} - H_{S,t-1}\right)}_{\text{net housing purchases}} - \underbrace{\delta p_t H_{S,t-1}}_{\text{maintenance}} + \underbrace{T_{S,t}}_{\text{rebates}}$$

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

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). In addition to the budget constraint, the saver must also satisfy the fixed housing demand constraint (friction) $H_{S,t} = \bar{H}_S$ at all times.

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:

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

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:

$$\max_{L_t,Z_t} p_t L_t^{\varphi} Z_t^{1-\varphi} - p_{\text{Land},t} L_t - Z_t$$

where $p_{\text{Land},t}$ is the equilibrium price of land permits, and the price of nondurables is normalized to unity.

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.

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

6.1 Model Solution

We now present the key equilibrium conditions of the model, while reserving the full set of equilibrium conditions to Appendix A. These key equations are the optimality conditions for borrower and landlord housing, which correspond to the inverted demand and supply curves

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

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

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

These equations map directly into the supply and demand framework of Section 2, where the slopes 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.

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.

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 influenced by credit conditions. Correspondingly, we identify a household as a "borrower" if it either (i) owns a home, and whose mortgage balance net of liquid assets is greater than 30% of the home's value, or (ii) does not own a home. We believe both of these groups would likely find it very difficult to purchase a home without credit. This procedure yields a population share of $\chi_B = 0.626$ and an income share of $s_B = 0.525$. For landlord demographics, we consider the limit $\chi_L \to 0$ and assume that landlords do not receive labor income, instead subsisting entirely on their rental earnings.

Turning to preferences, the key parameter is the borrower's discount factor, β_B, which determines how much latent demand for credit there is in the economy, and in turn, how

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	0.525	N	1998 SCF
Landlord pop. share	χ_L	0.000	N	Normalization
Borr. discount factor	β_B	0.974	Y	PMI Rate (see text)
Saver discount factor	β_S	0.992	Y	Nom. interest rate = 6.46%
Landlord discount factor	β_L	0.974	Y	Equal to β_B
Housing utility weight	ξ	0.200	N	Davis and Ortalo-Magné (2011)
Saver housing demand	$\bar{H}_S$	5.299	Y	Steady state optimum
Ownership Benefit Heterogeneity				
Landlord het. (location)	$\mu_{\omega,L}$	-0.002	Y	Avg. homeownership rate
Landlord het. (scale)	$\sigma_{\omega,L}$	0.020	Y	Empirical elasticities (Section 7)
Borr. het. (location)	$\mu_{\omega,B}$	0.004	Y	Borr. VTI (1998 SCF)
Borr. het. (scale)	$\sigma_{\omega,B}$	0.008	Y	Implied subsidy (see text)
Technology and Government				
New land per period	$\bar{L}$	0.090	Y	Residential inv = 5% of GDP
Land share of construction	φ	0.371	N	Res inv. elasticity in boom
Housing depreciation	δ	0.005	N	Standard
Inflation	$\bar{\pi}$	1.008	N	3.22% Annualized
Tax rate	τ	0.204	N	Standard
Mortgage Contracts				
Refinancing rate	$\bar{\rho}$	0.034	N	Greenwald (2018)
Loan amortization	ν	0.45%	N	Greenwald (2018)
Borr. LTV Limit	θ_B^{LTV}	0.850	N	Greenwald (2018)
Borr. PTI Limit	θ_B^{PTI}	0.360	N	Greenwald (2018)
Borr. PTI offset (taxes etc.)	α_B	0.09%	N	Greenwald (2018)
Landlord LTV Limit	θ_L^{LTV}	0.000	N	No landlord credit

much a relaxation of credit will influence house prices. We infer this parameter from the pricing on private mortgage insurance (PMI) — the additional fees and interest rates that a borrower must pay in order to obtain a high-LTV loan. This approach is motivated by the fact that many borrowers choose to pay for PMI, while many do not, meaning that the typical borrower should be close to indifferent.[16] We choose β_B so that the typical borrower would be indifferent between receiving a loan at 80% LTV, and paying the exact

[16] For example, 37.7% of Fannie Mae purchase loans required PMI over the 1999-2008 boom period (source Fannie Mae Single Family Dataset).

FHA insurance scheme for a loan at 95% LTV: an up front fee of 1.75% of the loan, plus a spread of 80bp.[17]

For the other preference parameters, we assume a standard consumption weight parameter of $\xi = 0.200$ on housing, following the evidence in Davis and Ortalo-Magné (2011). 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). We set the saver's fixed level of demand $\bar{H}_S$ equal to the level they would choose in steady state at prevailing prices. This implies that while saver demand is fixed in the short run, it is at the correct "long run" equilibrium value.

Lastly, we set the landlord discount rate β_L to be equal to β_B. This calibration is convenient for experiments involving overoptimistic house price expectations, which we impose as news about high expected future rents. When borrowers and landlords have different discount factors, whoever is more patient will value these future rental services or cashflows more, and will begin accumulating housing today, shifting the homeownership rate. To avoid taking a stand on how overoptimism influenced homeownership during the boom, we choose as our baseline the symmetric case where expectations influence both borrower and landlord valuations symmetrically.

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 borrower distribution, we set $\mu_{\omega,B}$, the typical ownership utility for borrowers, to target the average ratio of home value to income among "borrowers" who own homes in

[17] We choose the FHA scheme because it is much simpler to implement in the model than the GSE insurance scheme, where pricing is less transparent, and insurance premia are only paid until the borrower's LTV drops below 80%. However, the overall costs of the two forms of insurance are similar, as can be seen in e.g., Goodman and Kaul (2017).

the 1998 SCF, equal to 8.81 (quarterly). We found the ownership benefit dispersion parameter $\sigma_{\omega,B}$ difficult to pin down in the data, but thankfully this parameter plays little role in our results. To provide a sense of scale, we set $\sigma_{\omega,B}$ so that 5% of borrowers (10% of borrowers who own) would instead prefer to rent at steady state if given a 15% rental subsidy. For the landlord, we set $\mu_{\omega,L}$ to attain the correct homeownership rate among "borrowers" in the 1998 SCF (49.64%). Since all savers own, this ensures an overall homeownership rate of 68.50% at steady state, matching the 1998 SCF. The landlord dispersion 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.

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 function φ, we note that $\alpha/(1-\alpha)$ is the elasticity of residential investment to house prices, and choose 0.371 so that this elasticity is equal to the ratio of the peak log increase in the residential investment share of output to the peak log increase in prices over the boom. 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

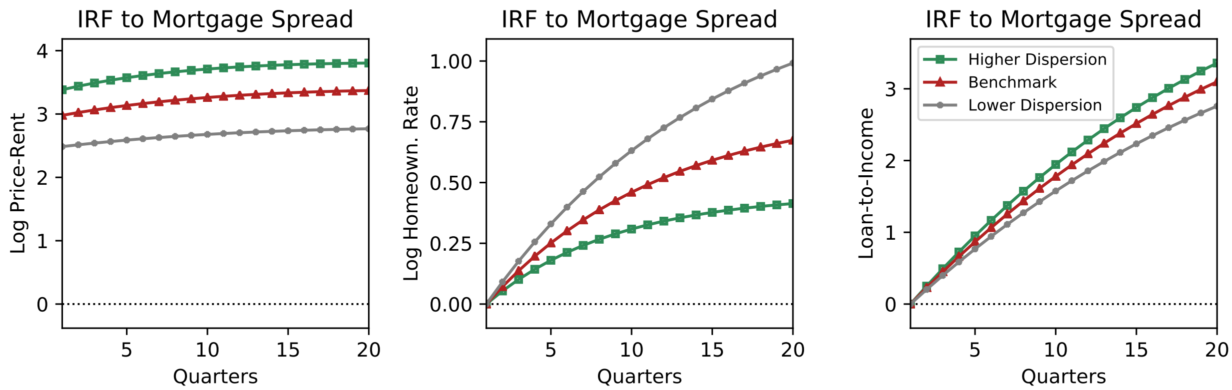

Figure 5: Model Response to our Identified Credit (CLL) Shock by $\sigma_{\omega,L}$

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 5 shows this impulse response for three possible calibrations of our dispersion parameter, $\sigma_{\omega,L}$. Our benchmark calibration of $\sigma_{\omega,L} = 0.020$ 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 3.37 and 0.67, respectively. The other two series show paths for alternative calibration choices: the "Higher Dispersion" series doubles the dispersion to $\sigma_{\omega,L} = 0.030$, while the "Lower Dispersion" series halves the dispersion to $\sigma_{\omega,L} = 0.007$. 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 9.3 (3.80 vs. 0.41, respectively) while the "Lower Dispersion" series generates a 20Q ratio of 2.8 (2.77 vs. 0.99, respectively). Overall, this comparison illustrates that $\sigma_{\omega,L}$ can be identified relatively narrowly using this procedure.

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

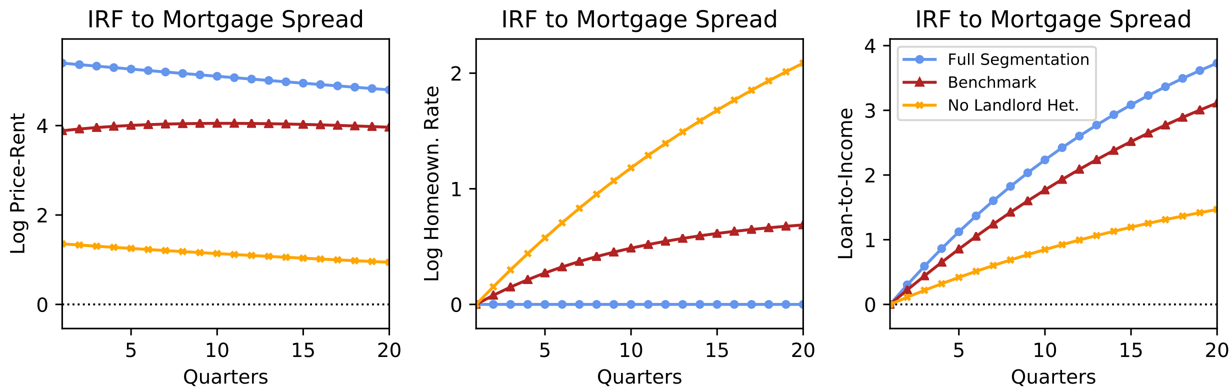

Figure 6: Model Response to our Identified Credit Shock, Benchmark vs. Polar Cases

markets ($\sigma_{\omega,L} \to 0$).

To this end, Figure 6 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, and displaying less than half of the increase in the homeownership rate (0.69% vs. 2.09%), 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 6. 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 twice as much in the Benchmark vs. No Landlord Heterogeneity economy (3.11% vs. 1.47%), 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 83% of the 20Q rise in both price-to-rent and loan-to-income 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.

As an additional check against our empirical results and the existing literature, we consider the effects on prices and rents separately in Appendix Figure B.1. These results indicate that our benchmark price-rent response is driven by a large increase in house prices, with an elasticity around 3.4, and a much smaller increase in rents, consistent (up to scale) with our empirical findings. This elasticity of house prices to interest rates is consistent with estimates in the literature, for instance Adelino et al. (2012) who find that this elasticity measured from a similar shock to the conforming loan limit lies in the interval $[1.2, 9.1]$. The full segmentation model also delivers an elasticity in this range, while a model with no heterogeneity implies a house price elasticity close to zero, consistent with our theory in Section 2 but at odds with this empirical finding.

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 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 7. To highlight the role of landlord heterogeneity, we again plot the responses in our Benchmark model against those of

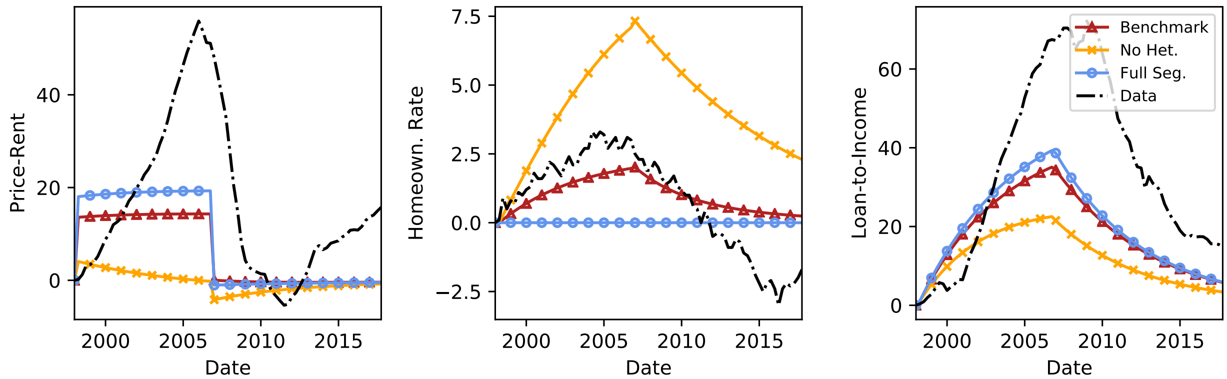

Figure 7: Credit Relaxation Experiment

alternative models with no landlord heterogeneity and full segmentation (infinite 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 28% 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 dynamics in the Benchmark model are much closer to the "Full Segmentation" model, where this credit relaxation would account for 38% 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 7, credit growth over the boom is much larger in the Benchmark economy vs. the No Heterogeneity economy, explaining 52% vs. 33% of the observed rise. This additional credit growth is a direct consequence of the larger house price appreciation in the Benchmark 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

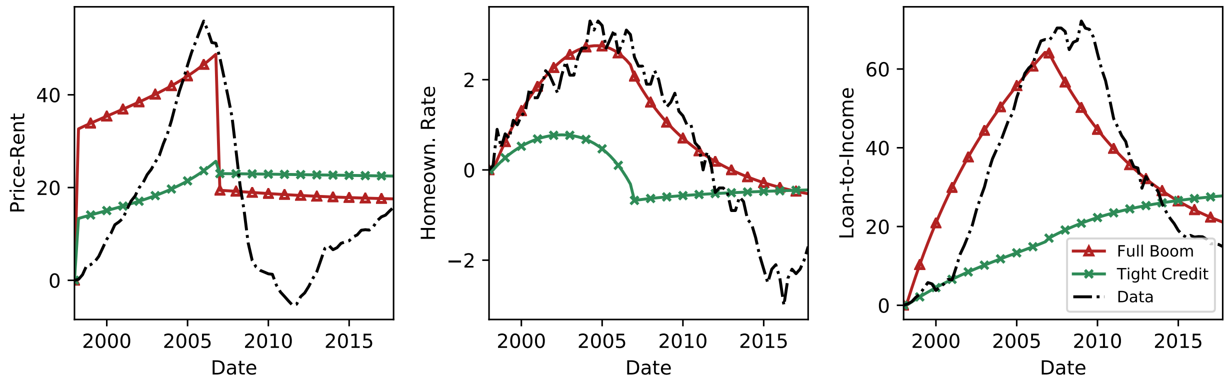

Figure 8: Complete Boom Experiment

price rise, and therefore delivers credit growth corresponding to 58% 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, shown in Appendix Figure B.2 can explain 60% of the rise in price-to-rent ratios and 80% 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 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.[18] 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 agents believe in a future increase in ξ, the expenditure share on housing services, of 30%. Given the size of this required expected growth, we note that part if not all of this supply shift

[18]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.

required to fit the data could alternatively be generated by shifts in landlord access to credit, which we discuss in Section 9.[19]

To match the bust, we impose a further 3ppt fall in interest rates, which is also reflected in a 3ppt 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.[20] 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 financial 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 "Tight Credit" 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 47% and in loan-to-income ratios by 75% relative to our Benchmark scenario. These contributions, which provide the upper bound for our estimated role of credit during the boom-bust, are larger than the shares explained by relaxing credit in isolation (Figure 7), because in the Benchmark model, loose credit also amplifies the other components of the boom, particularly the role of house price expectations. The simple intuition, discussed at length in Greenwald (2018), is that even if households are very optimistic about future cash flows from housing, binding PTI limits limit households' ability to finance these properties in the absence of a corresponding rise in income, dampening prices and credit growth when credit remains tight.

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,

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

[20] 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.

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.

9 Model Extensions

In the model as presented so far, the only credit insensitive agents who could enter the owner-occupied market are deep-pocketed landlords who face heterogenous costs in converting properties between owner-occupied and renter-occupied. While this makes the economics of the model clear, it may not be realistic. In particular, it may be that these landlords also face financial constraints. Furthermore, it may also be the case that unconstrained savers can trade housing with borrowers and landlords. In this section, we extend the model to include these two forces, first considering landlords credit and then saver housing demand.

9.1 Landlord Credit

We begin by introducing landlord credit. In practice, landlords are not deep-pocketed, and nearly every investor-owned property is purchased with a mortgage.

Supply and Demand Intuition. We begin by reassessing the intuition developed in Section 2. Recall we developed a supply-and-demand model in homeownership rate vs. price-to-rent ratio space. The demand curve represents demand by constrained homeowners, while the supply curve represents supply from landlords who convert renter housing to owner-occupied housing and sell it to households. A credit relaxation shifts the demand curve, which leads to movement in the price-to-rent ratio with no homeownership rate movement if rental and ownership are perfectly segmented (Figure 2a) and movement in the homeownership rate with no price-to-rent ratio movement if conversion is perfectly frictionless and (Figure 2b). With limited conversion, the supply curve is upward sloping and there is movement in both the homeownership rate and the price-to-rent ratio (Figure 2c).

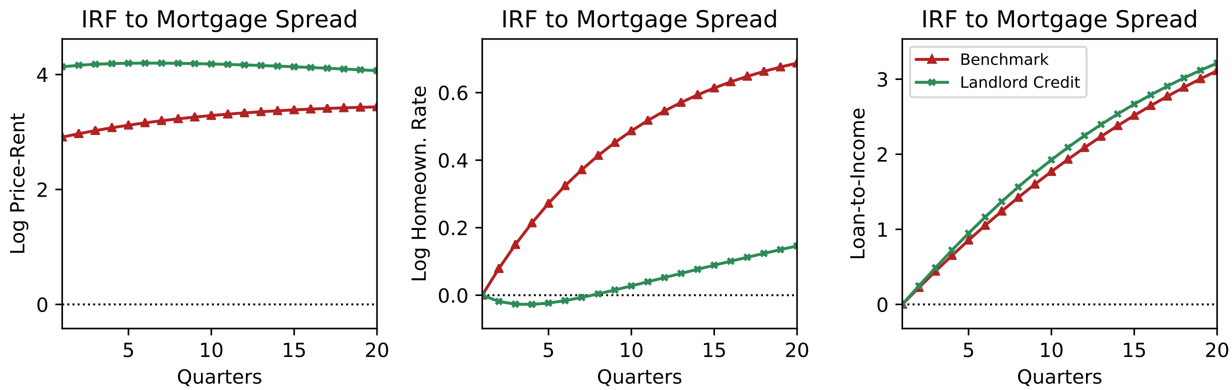

Figure 9: Effect of Landlord Credit on Model Response to Credit (CLL) Shock

Introducing credit for landlords implies that a relaxation in credit not only shifts the demand curve upward but also shifts the supply curve upward as in Figure 2d. The degree of the supply curve shift is endogenous and depends on the model's parameters. Due to this shift, adding landlord credit to the baseline model while holding the parameters fixed would lead to a smaller change in the homeownership rate (and potentially a *negative* change if the supply curve shifts enough) and a larger change in the price-to-rent ratio, as represented by a shift from the solid supply curve to the dashed supply curve in Figure 2d.

To illustrate this intuition quantitatively (while still holding parameters fixed), we implement a version of the model with landlord credit. Specifically, we assume that landlords face an LTV limit of 65%, – the standard constraint for multi-family construction loans – and do not face a PTI limit.[21] Figure 9 compares the responses to our same credit shock of this model with our benchmark calibration with no landlord credit assuming that the shock causes mortgage spreads for households and landlords fall equally. In terms of magnitudes, the shift in the supply curve due to landlord credit modestly increases the response of the price-rent ratio (3.44% to 4.07%) at the 20Q horizon, while cutting the rise in the homeownership rate close to zero (0.69% to 0.15%) over the same period.

There are three important takeaways from this plot. The first is that the presence of landlord credit can act as a quantitatively important substitute for rental frictions, with the Landlord Credit line in Figure 9 looking very similar to the Full Segmentation re-

[21] These assumptions correspond to the parameter values $\theta_L^{LTV} = 0.65$ and $\theta_L^{PTI} = \infty$, and imply $F_{L,t}^{LTV} = 1$.

sponse in Figure 6, despite having the same rental friction as the Benchmark model. Second, landlord reliance of credit can potentially generate a zero or nonpositive response of homeownership to credit shocks, which cannot be rejected by our empirical estimates, and may explain our finding of nonpositive point estimates in Figure 3. Finally, to transparently show the effect of the supply shift, we have not recalibrated the model to model to match our target moments when we added landlord credit. Matching target price-rent relative to homeownership causal effect would imply a lower level of rental frictions in the landlord credit model, leading to highly similar dynamics across the two specifications.

9.2 Saver Housing Demand

A key assumption in our modeling framework is that savers own and consume a fixed amount of 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, as in Kiyotaki et al. (2011). 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, potentially allowing savers to buy up units of housing and incorporate 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. It is important to note that this margin is also typically available to agents even in heterogeneous agent models with a discrete set of house sizes. Although households must choose among these fixed sizes, the housing stock can still be costlessly reallocated among the different sizes, for instance to allow several smaller houses to transform into one larger one.

This margin is likely much more restricted in reality, as housing in the real world oc-

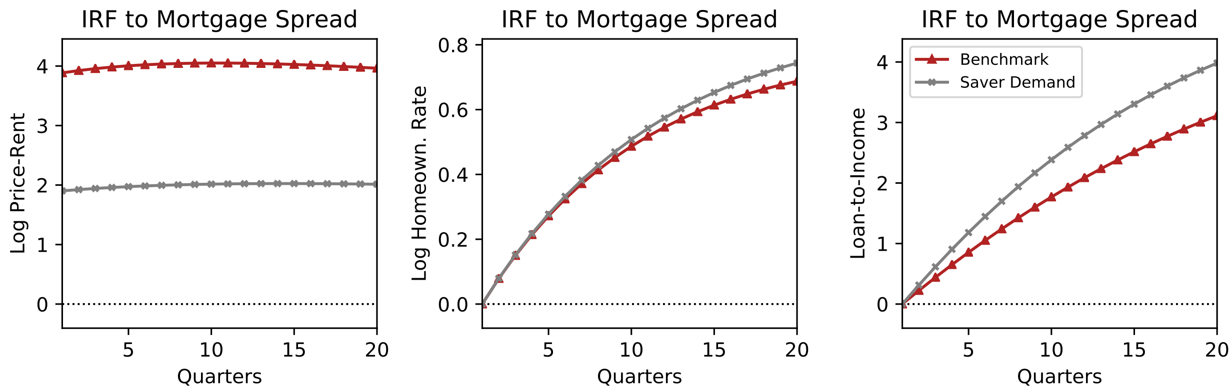

Figure 10: Effect of Landlord Credit on Model Response to Credit (CLL) Shock

cupies vast spectra of quality levels and geographic locations, implying that households cannot easily enjoy more service flows by purchasing additional housing. Landvoigt et al. (2015) formalize this argument by endogenizing segmentation in a sorting model and show that credit has significant effects on the price of lower-quality houses, and we see our assumption of segmentation as a reduced-form version of what they show endogenously. Furthermore, 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.

We once again demonstrate this with a quantitative example. To do this, we extend the baseline model to allow the saver households to flexibly transact with borrowers and landlords in the housing market. This implies the additional optimality condition from the saver first order condition:

$$p_t^{\text{Saver}} = \mathbb{E}_t\Big\{\Lambda_{t+1}^S\Big[\underbrace{u_{h,t}^S/u_{c,t}^S}_{\text{housing services}} + \underbrace{\big(1-\delta\big)p_{t+1}}_{\text{continuation value}}\Big]\Big\},$$

where p_t^{Saver} must equal the market house price at equilibrium. To clear this market, the savers adjust their housing stock $H_{S,t} = \bar{H}_t - H_{B,t} - H_{L,t}$ by transacting in housing, which in turn adjusts the marginal rate of substitution $u_{h,t}^S/u_{c,t}^S$.

Figure 10 shows the response to our credit shock in this extension, denoted "Saver Demand," alongside our Benchmark results. As can be seen, adding a flexible saver margin substantially decreases the response of the price-rent ratio by nearly half (3.96% to 2.01%) while leaving the homeownership rate unchanged (0.69% vs. 0.74%) at the 20Q horizon. The dampening effect of saver demand on prices is quantitatively large, and could be even stronger for saver preferences with less curvature.[22] By contrast, saver demand matters little for the homeownership rate because model savers always own and never rent, with their transactions affecting the intensive margin of house size. As a result, these transactions have no direct effect on the overall homeownership rate.

Since we are skeptical that unconstrained households buying and selling second homes was an important factor during the boom-bust, we have chosen to shut down this margin in our benchmark model. Nonetheless, even if this margin was active, we believe our calibration procedure ensures that our results should not be seriously biased by this assumption. 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. As shown above, if unconstrained savers are indeed transacting with borrowers in the data, this should dampen the house price response to credit, while leaving the homeownership rate response unchanged. Our calibration procedure would therefore require even more extreme rental market frictions to reproduce our empirical findings, restoring the impact of credit on prices, and leading to similar overall results.

10 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 most of the literature has focused on two polar cases with regards to the segmentation in housing markets between credit-insensitive agents such as landlords and unconstrained savers and credit-sensitive agents.In the first, landlords are unable to convert homes from owner-occupied to renter-

[22]The degree of dampening is determined by how strongly the marginal rate of substitution term $u^S_{h,t}/u^S_{c,t}$ varies with $H_{S,t}$. For example, with linear saver preferences $u^S(c,h) = c + \xi h$, we would obtain the constant value $u^S_{h,t}/u^S_{c,t} = \xi$, implying that house prices would be completely invariant to a credit shock.

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

ADELINO, M., A. SCHOAR, AND F. SEVERINO (2012): "Credit Supply and House Prices: Evidence From Mortgage Market Segmentation," Tech. rep., National Bureau of Economic Research.

AGARWAL, S., G. AMROMIN, I. BEN-DAVID, S. CHOMSISENGPHET, T. PISKORSKI, AND A. SERU (2017a): "Policy Intervention in Debt Renegotiation: Evidence From the Home Affordable Modification Program," *Journal of Political Economy*, 125, 654–712.

AGARWAL, S., G. AMROMIN, S. CHOMSISENGPHET, T. LANDVOIGT, T. PISKORSKI, A. SERU, AND V. YAO (2017b): "Mortgage Refinancing, Consumer Spending, and Competition: Evidence from the Home Affordable Refinancing Program," Working Paper 21512, National Bureau of Economic Research.

CAMPBELL, J. Y., N. CLARA, AND J. F. COCCO (2018): "Structuring Mortgages for Macroeconomic Stability," Working Paper, Harvard University and London Business School.

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

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

DI MAGGIO, M. AND A. KERMANI (2017): "Credit-Induced Boom and Bust," *The Review of Financial Studies*, 30, 3711–3758.

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.

GLAESER, E. L., J. D. GOTTLIEB, AND J. GYOURKO (2012): "Can Cheap Credit Explain the Housing Boom?" in *Housing and the Financial Crisis*, University of Chicago Press, 301–359.

GOODMAN, L. AND K. KAUL (2017): "Sixty Years of Private Mortgage Insurance in the United States," *Urban Institute, August*.

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

GREENWALD, D., T. LANDVOIGT, AND S. VAN NIEUWERBURGH (2017): "Financial Fragility With SAM?" MIT Sloan Research Paper No. 5261-17.

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.

HEDLUND, A. (2016): "The Cyclical Dynamics of Illiquid Housing, Debt, and Foreclosures," *Quantitative Economics*, 7, 289–328.

HUO, Z. AND J.-V. RIOS-RULL (2016): "Financial Frictions, Asset Prices, and the Great Recession," Working Paper, Federal Reserve Bank of Minneapolis.

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.

——— (2017): "The Mortgage Rate Conundrum," Working Paper.

KAPLAN, G., K. MITMAN, AND G. L. VIOLANTE (2019): "The Housing Boom and Bust: Model Meets Evidence," Tech. rep., National Bureau of Economic Research.

KIYOTAKI, N., A. MICHAELIDES, AND K. NIKOLOV (2011): "Winners and Losers in Housing Markets," *Journal of Money, Credit and Banking*, 43, 255–296.

LANDVOIGT, T., M. PIAZZESI, AND M. SCHNEIDER (2015): "The Housing Market(s) of San Diego," *American Economic Review*, 105, 1371–1407.

LIU, Z., P. WANG, AND T. ZHA (2019): "A Theory of Housing Demand Shocks," Working Paper 25667, National Bureau of Economic Research.

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

MIAN, A. AND A. SUFI (2019): "Credit Supply and Housing Speculation," Working Paper 24823, National Bureau of Economic Research.

MITMAN, K. (2016): "Macroeconomic Effects of Bankruptcy and Foreclosure Policies," *American Economic Review*, 106, 2219–2255.

ORTALO-MAGNÉ, F. AND S. RADY (2006): "Housing Market Dynamics: On the Contribution of Income Shocks and Credit Constraints," *Review of Economic Studies*, 73, 459–485.

PISKORSKI, T. AND A. TCHISTYI (2017): "An Equilibrium Model of Housing and Mortgage Markets With State-Contingent Lending Contracts," Working Paper, Columbia University and University of Illinois.

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.

A Equilibrium Conditions

This section presents the full set of equilibrium conditions of the model.

Borrower's Problem. The borrower's optimality conditions are:

$$(h_{B,t}): \qquad q_t = (u^h_{B,t}/u^c_{B,t}) \tag{6}$$

$$(H^*_{B,t}): \qquad p_t = \frac{\mathbb{E}_t\left\{\Lambda_{B,t+1}\left[\bar{\omega}_{B,t} + q_{t+1} + \Big(1-\delta-(1-\rho_{B,t+1})\mathcal{C}_{B,t+1}\Big)p_{t+1}\right]\right\}}{1-\mathcal{C}_{B,t}} \tag{7}$$

$$(M^*_{B,t}): \qquad 1 = \Omega^B_{M,t} + r^*_{j,t}\Omega^B_{X,t} + \mu_{B,t} \tag{8}$$

$$(\rho_{B,t}): \qquad \rho_{B,t} = \Gamma_\kappa\Bigg\{\underbrace{(1-\Omega^B_{M,t}-\Omega^B_{X,t}\bar{r}_{B,t-1})\left(1-\frac{(1-\nu)\pi_t^{-1}M_{B,t-1}}{M^*_{B,t}}\right)}_{\text{new debt incentive}} - \underbrace{\Omega^B_{X,t}\left(r^*_{B,t}-\bar{r}_{B,t-1}\right)}_{\text{interest rate incentive}}\Bigg\}, \tag{9}$$

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:

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

Equation (6) sets the rent equal to the marginal rate of substitution between housing services and consumption. Equation (7) 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 (8) 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 (9) 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 optimality conditions are

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

where $\mathcal{C}_{L,t} = \mu_{L,t}F^{LTV}_{L,t}\theta^{LTV}_{L,t}$ is defined analogously to the borrower case. The fixed point conditions that pin down the marginal continuation costs of debt are defined by

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

symmetric to the borrower case.

Saver's Problem. The saver's optimality conditions are:

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

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

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

Construction Firm's Problem. The construction firm's optimality conditions are

$$p_{\mathrm{Land},t} = p_t \varphi L_t^{\varphi-1} Z_t^{1-\varphi}$$
$$1 = p_t (1-\varphi) L_t^{\varphi} Z_t^{-\varphi}.$$

B Additional Figures

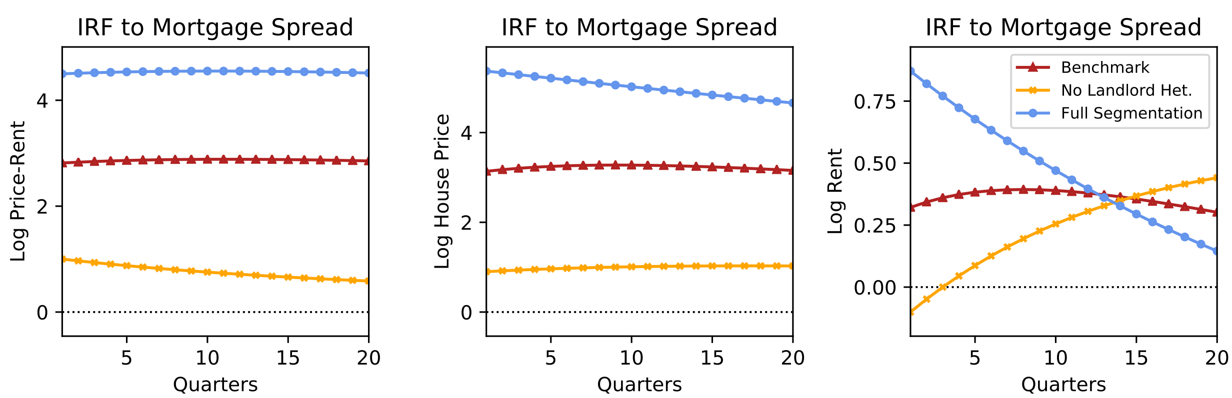

Figure B.1: Model Response to our Identified Credit (CLL) Shock by $\sigma_{\omega,L}$

C Empirical Appendix

This section presents additional empirical results and robustness checks.

C.1 State-Level Results

As mentioned in section 3, our measure of the homeownership rate at the CBSA level is complicated by changing CBSA definitions. We drop periods in which CBSA defini-

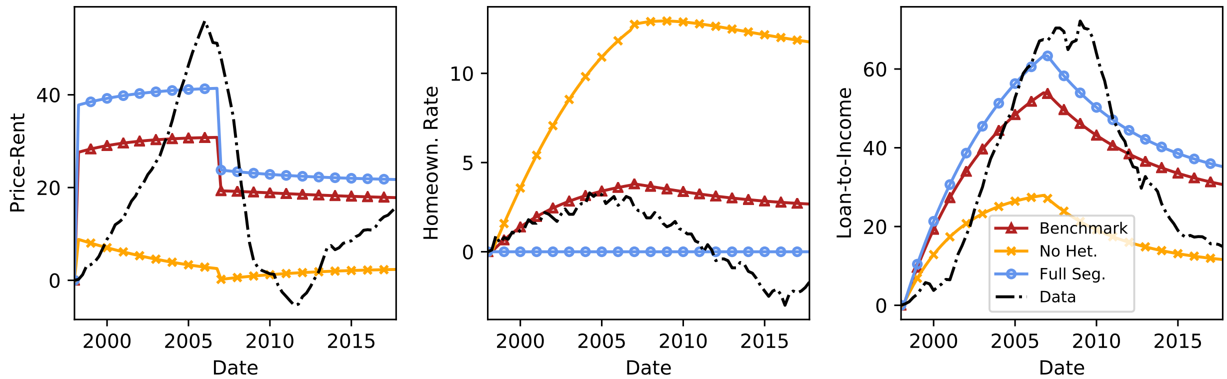

Figure B.2: Credit Relaxation + Fall in Interest Rates

tions change, but this provides an unbalanced panel and may bias the results. To address this concern, in this section we repeat our analysis for a state-level panel. Because state definitions do not change, the homeownership rate panel is of higher quality.

We create a state-level data set that includes house prices, homeownership rates, and measures of credit availability. For house prices, we use the FHFA annual house price indices. For homeownership rates, we use data from the Census. For credit and the Loutskina-Strahan instrument, we use HMDA and collapse the microdata as described in the data appendix, except at the state level instead of the CBSA level. Unfortunately, we are unable to obtain the high-quality rent data from CBRE at the state level, and so our analysis does not include rents or price-to-rent ratios. We also need a state-level housing supply elasticity. For this, we use a weighted average of the Saiz (2010) elasticity at the MSA level for the MSAs in each state using 2000 population as the weights. We then use the exact same econometric approach as in Section 4.

Unfortunately, our instrument is far less powerful than before. The first state F statistic ranges from 1.5 for the number of loans to 5 for the volume of loans.[23] This makes sense: At higher levels of aggregation, the sharp micro variation is averaged out leading to a weaker instrument. Consequently, our results are much more imprecise than with our main instrument. Given this, we consider an alternate instrument as well which uses the national conforming loan limit to determine the fraction near the conforming loan limit

[23] Given the weak instrument, we ensure these results are robust to using LIML instead of two stage least squares but report two-stage least squares results below.

rather than the local conforming loan limit. This instrument is marginally more powerful, and we see it largely as a robustness check.

Table C.1: State-Level IV Results For Contemporaneous Outcomes

	$\Delta\log(\text{Price})$			$\Delta\log(\text{Homeownership Rate})$		
$\Delta\log(\#\text{ Loans})$	0.254			0.008		
	(0.345)			(0.093)		
$\Delta\log(\text{Vol. Loans})$		0.124			0.005	
		(0.117)			(0.053)	
$\Delta\log(\text{Loan/Income})$			0.455			0.022
			(0.316)			(0.247)
N	1274	1274	1274	1274	1274	1274
	$\Delta\log(\text{Price})$ (Alt Instrument)			$\Delta\log(\text{Rent})$ (Alt Instrument)		
$\Delta\log(\#\text{ Loans})$	0.290			0.015		
	(0.302)			(0.069)		
$\Delta\log(\text{Vol. Loans})$		0.152			0.007	
		(0.114)			(0.040)	
$\Delta\log(\text{Loan/Income})$			0.604			0.045
			(0.316)			(0.200)
N	1274	1274	1274	1274	1274	1274

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.

Table C.1 shows the panel IV results, and Figure C.3 shows the local projection impulse responses. The contemporaneous response of prices and the impulse response for prices are broadly consistent with the CBSA results both qualitatively and quantitatively, albeit with much larger standard errors. Nonetheless, the fact that the point estimates are so similar is reassuming and confirms that our main results for price seem to hold in our state panel.

By contrast, the response of homeownership rates has point estimates near zero, although we cannot statistically reject equality with the price coefficients given the large standard errors on price. Similarly, the homeownership rate insignificant and negative at all horizons. While this does not provide strong statistical evidence on the response of homeownership rates to credit supply shocks, it does not seem to indicate that data issues

Figure C.3: State-Level Panel Local Projection IV Impulse Responses

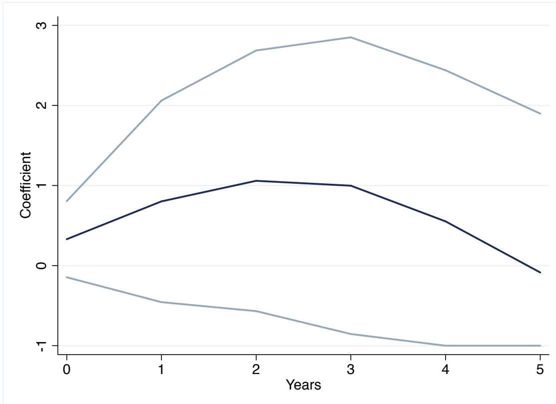

(a) Price

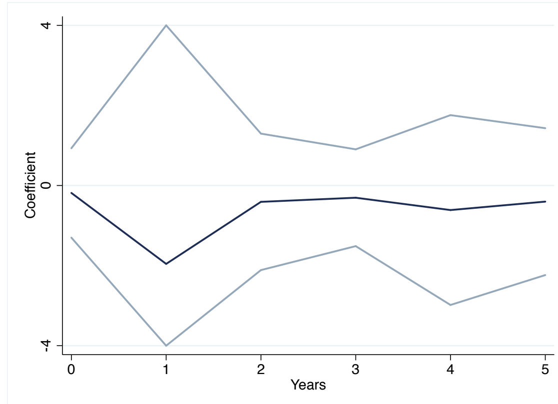

(b) Homeownership Rate

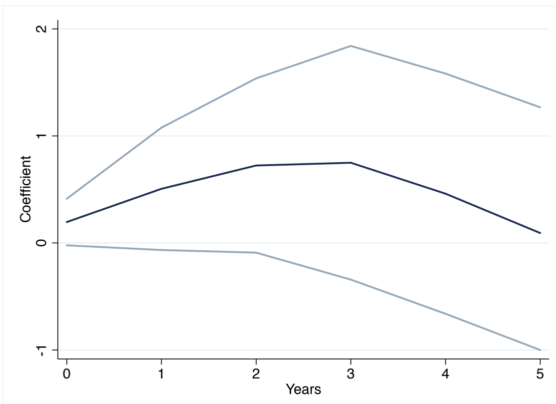

(c) Price (Alternate Instrument)

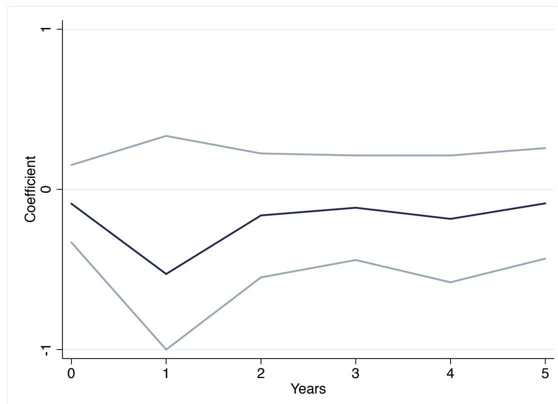

(d) Homeownership Rate (Alternate Instrument)

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 $Fraction_{i,t-1}$, $Fraction_{i,t-2}$, $Fraction_{i,t-2} \times \%\Delta CLL_t - 1$, and two lags of the outcome variable. The alternate instrument computes $Fraction$ using the national conforming loan limit rather than the local conforming loan limits for each county, which can be higher than the national. All standard errors are robust.

with the CBSA-level homeownership rates explain our results.

Overall, we see the state level analysis as providing some confirmation – albeit statistically insignificant – that our main results are not driven by data quality issues.