

Do Credit Conditions Move House Prices?

Daniel L. Greenwald* and Adam Guren[‡]

November 2021

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 that 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 elasticities of the price-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 three different credit supply instruments and use it to calibrate our model. Our results reveal that rental markets are highly frictional and closer to fully segmented, which implies large effects of credit on house prices. In particular, changes to credit standards can explain between 34% and 55% 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, Nelson Lind, and seminar participants at the 2019 AEA Annual Meeting, Duke University, SED, SITE, NYU, the Federal Reserve Bank of Chicago, the University of Nebraska, the Bank of Ireland, and VMACS 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 to the effectiveness of macroprudential policy tools, yet more than a decade later 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. In contrast, papers such as Kaplan, Mitman, and Violante (2020) argue that credit conditions explain virtually none of the boom and bust in house prices, which are instead dominated by changes in beliefs.¹

This paper makes sense of these divergent findings by elucidating the source of disparate results and quantitatively assessing the role of credit in driving the 2000s housing boom and bust.² Our analysis proceeds in four steps. First, we illustrate why existing models are at odds. Second, we develop and implement an empirical strategy to estimate where reality falls on the spectrum of possible model assumptions. Third, we construct a modeling framework flexible enough to nest this spectrum and use our empirical findings to pin down where the economy lies on that spectrum. Fourth, we use our calibrated model to quantify the role of credit changes in driving the 2000s housing boom and bust.

To begin, we show that the key difference between these disparate findings is the degree to which credit insensitive agents can absorb credit-driven demand by constrained agents, which in turn depends on the degree of segmentation in housing markets. This mechanism is clearest and most important in the rental market. In models with perfectly segmented rental markets (most commonly, these assume no rental market exists), favorable credit conditions increase demand for housing by borrowers who compete with each other for the same properties, bidding up house prices. In contrast, models with no rental market segmentation feature deep-pocketed landlords who are willing to trade an unlimited amount of housing at a price equal to the present value of rents. Since landlords are

¹For more examples, Landvoigt, Piazzesi, and Schneider (2015), Greenwald (2018), Guren, Krishnamurthy, and McQuade (2021), Garriga, Manuelli, and Peralta-Alva (2019), Garriga and Hedlund (2020), Garriga and Hedlund (2018), Justiniano, Primiceri, and Tambalotti (2019), and Liu, Wang, and Zha (2019) analyze models that imply credit conditions played a key role in the boom and bust, while Kiyotaki, Michaelides, and Nikolov (2011) study a model in which credit conditions played only a limited role.

²Since credit standards are endogenously set by lenders, the division between credit factors and other factors may not be obvious. For example, overoptimistic beliefs can both raise house prices on their own, and cause lenders to relax credit standards (Foote, Gerardi, and Willen (2012)). In this paper, we define the role of credit to be the difference in outcomes between what occurs when credit conditions change compared to a counterfactual in which credit conditions were exogenously held fixed, regardless of the ultimate cause of the shift in credit conditions.

assumed not to use credit and credit has little effect on rents, changes in credit conditions do not influence landlords' reservation price, and the price is effectively fixed. As a result, a credit expansion leads many households to buy housing from landlords, increasing the homeownership rate but not house prices.

The vast majority of models in the literature fall under one of these two paradigms, which can be represented by perfectly inelastic and perfectly elastic "tenure supply" curves that represent the relative price schedules at which landlords are willing to sell rental housing to households. In this paper, we instead allow for the possibility of intermediate frictions between these two extremes, with the relative strength of the price-rent and homeownership responses determined by the slope of the tenure supply curve. This slope consequently provides a new and important empirical moment to be matched by any model seeking to study the influence of credit on house prices.

To measure the slope of the tenure supply curve, we follow the traditional approach of using demand instruments. We estimate this slope as the ratio of the elasticity of the price-rent ratio to an identified credit shock, compared to the elasticity of the homeownership rate to that same shock. For our identified credit shock, we look to the existing literature for three different identification strategies. The first and most statistically powerful approach follows Loutskina and Strahan (2015) (hereafter LS) in instrumenting for local credit using differential city-level exposure to changes in the conforming loan limit. The second approach follows Di Maggio and Kermani (2017) (hereafter DK) in exploiting the 2004 preemption of state-level anti-predatory-lending laws for national banks by the Office of the Comptroller of the Currency. The third and final approach follows Mian and Sufi (2021) (hereafter MS) who follow Nadauld and Sherlund (2013) in using differential city-level exposure to a rapid expansion of the private label securitization market through heterogeneity in bank funding sources.

Despite relying on different sources of identification and operating through different segments of the mortgage market, our three empirical approaches deliver largely consistent findings. All three sets of approaches indicate that shocks to credit supply significantly increase house prices and the price-rent ratio, but have much smaller and statistically insignificant effects on the homeownership rate.³ The resulting slope estimates range from three to infinity depending on the instrument and horizon, implying a substantial degree of segmentation.

To interpret these slopes economically, 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 het-

³Note that a shock to credit supply is a shock to housing *demand*, in line with our empirical strategy.

erogeneity in landlord and borrower preferences for ownership, which allows our model to feature a fractional and time-varying homeownership rate. Our framework nests both full segmentation and zero segmentation between rental and owner-occupied housing, as well as a continuum of intermediate cases. We calibrate the key parameter determining segmentation to directly match our empirical impulse responses, then use the model to compute the role of credit in driving the 2000s housing boom.

We find that a relaxation of credit standards in isolation explains 34% of the rise in the price-rent ratio observed in the boom, with a lower bound of 26% accounting for parameter uncertainty. These results contrast to -1% explained by the model with no segmentation, and the 38% explained by the model with full segmentation, implying that our estimated frictions are strong and closer to the fully segmented case. The role played by credit is even larger after incorporating changes in the price of credit. Combining a 2ppt decline in mortgage rates alongside the relaxation of credit standards allows our benchmark model to explain 72% of the observed rise in price-rent ratios, compared to 4% under no segmentation, and 82% under full segmentation.

Last, we consider experiments in which we add sufficient exogenous borrower and landlord demand shocks — for instance due to expectations of future prices or rents — to explain the entire boom in price-rent ratios and homeownership. Removing credit relaxation from this “full boom” scenario reduces the observed rise in price-rent ratios by 55%, which is larger than the 34% explained by credit standards in isolation due to strong nonlinear interactions between credit standards and non-credit factors. In contrast, a model with no segmentation calibrated to deliver the exact same aggregate rise in price-rent ratios and homeownership rates implies that removing credit relaxation reduces the rise in price-rent ratios only negligibly (5%). These sharply divergent results from models matching the same aggregate targets point to the importance of our empirical calibration approach. Our results imply that policies that restrict mortgage credit standards such as loan-to-value (LTV) and payment-to-income (PTI) ratios can slow house price growth, but only in the presence of the significant segmentation we find in our benchmark economy.

Our baseline model assumes that landlords do not use credit and that the saver housing stock is entirely segmented from the borrower housing stock. To close our analysis, we relax each of these assumptions in turn. When landlords use credit, a credit supply expansion also shifts supply outward. This leads to a larger price-rent ratio response and smaller homeownership rate response than under our benchmark model, implying our baseline results are a lower bound for the effect of credit on the price-rent ratio.

Finally, unconstrained savers can also dampen credit-induced house price fluctuations if their housing demand is not segmented from borrowers. We extend our model to allow

for frictionless trade between savers and borrowers and find that while this reduces the effect of credit on price-rent ratios by roughly 25%, changes in the price and quantity of credit still explain 54% of the observed rise in price-rent ratios over the boom period. Since this extension abstracts from important frictions on borrower-saver trade due to indivisibility and variation in the quality and location of housing, we consider it to be an extreme lower bound on the influence of credit on house prices.⁴

In summary, the ability of owners who do not use credit, such as landlords or unconstrained households, to absorb credit-driven demand by constrained households determines the extent to which shifts in credit supply influence house prices and the home-ownership rate. Our empirical finding that price-rent ratios respond significantly more to a credit supply shock than homeownership rates implies that this margin is subject to substantial frictions, so that prices do respond to credit in a meaningful way. Mapped into our structural model, these frictions are sufficient for credit changes to have driven an important share of the rise in house prices during the 2000s housing boom.

Related Literature. 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 (2020), Favara and Imbs (2015), Loutskina and Strahan (2015), Di Maggio and Kermani (2017), and Mian and Sufi (2021). These results, however, cannot be directly mapped into the share of the boom and bust explained by credit, since the quasi-random variation they use does not match the specific set of shocks that drove the boom and bust. We contribute to this literature by adding measures of the causal effect of credit on the homeownership rate and showing that the ratio of the responses of the price-rent ratio to those of the homeownership rate can identify structural elasticities. These elasticities can be mapped into a structural model to assess the effect of credit on house prices for an arbitrary set of shocks, including those that correspond to the 2000s boom and bust.

Given this focus, our closest empirical counterpart is Gete and Reher (2018), who also measure the impact of an identified credit shock on both the price-rent ratio and homeownership. Based on a different identification scheme than the three we use in this paper, Gete and Reher (2018) estimate a response of the price-rent ratio that is 85 times larger than the response of the homeownership rate. These estimates correspond to very strong frictions consistent with our estimates in Section 4.

⁴These absences are not specific to our framework but are nearly ubiquitous in the macro-housing modeling literature, which typically allows the total housing stock to be frictionlessly reshuffled between various houses and is thus inconsistent with indivisibility.

Our work also relates to a literature using quantitative models to study the effect of credit supply on house prices such as Favilukis et al. (2017), Kaplan et al. (2020), Kiyotaki et al. (2011), Greenwald (2018), Guren et al. (2021), Garriga et al. (2019), Garriga and Hedlund (2020), Garriga and Hedlund (2018), Justiniano et al. (2019), Liu et al. (2019), and Huo and Rios-Rull (2016). In this paper, we help explain the wide variation in results found by these models to date, and provide a new empirical framework to reconcile them.

Closest to our structure 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. Landvoigt et al. (2015) find that credit is important in explaining the larger boom observed at the bottom of the quality distribution. We see these results as highly complementary to our work. While our model of borrower-saver segmentation is much more simplistic, our tractable approach allows us to embed a similar set of frictions in a complete general equilibrium model of housing and mortgages that provides for a richer set of counterfactuals.

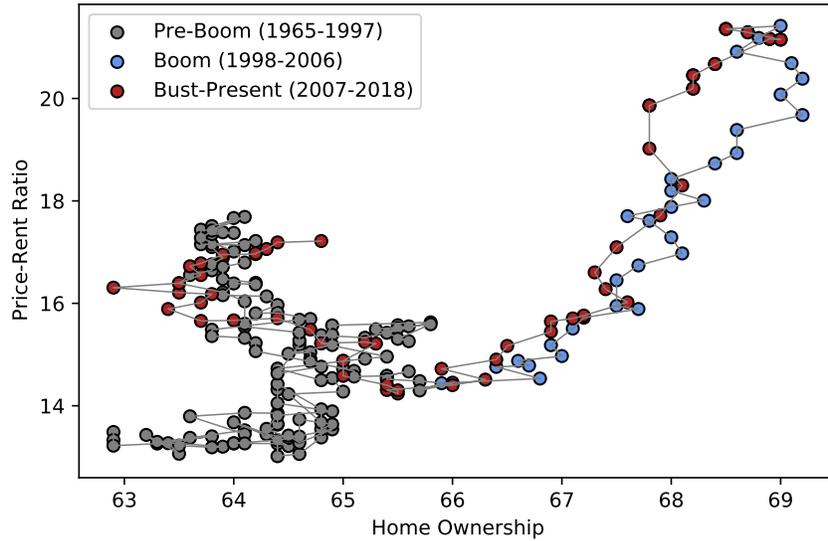
The modeling framework we employ also connects to work using tractable macro-housing models, including Campbell and Hercowitz (2005), Eggertsson and Krugman (2012), Garriga, Kydland, and Šustek (2017), Ghent (2012), Kiyotaki et al. (2011), Iacoviello (2005), Iacoviello and Neri (2010), Liu, Wang, and Zha (2013), Monacelli (2008), and Rog-nlie, Shleifer, and Simsek (2018). Our contribution relative to this literature is to provide a tractable methodology for incorporating fractional and time-varying homeownership rates and providing a new empirical moment to discipline it.

Overview. The rest of the paper is structured as follows. Section 2 presents the supply and demand model diagrammatically to provide intuition and motivate our estimation approach. Section 3 describes our data and empirical methodology. Section 4 presents our empirical results. Section 5 presents the model, Section 6 describes its calibration, and Section 7 presents our model results. Section 8 extends the model to include landlord credit and flexible saver housing demand. Section 9 concludes.

2 Intuition: Supply and Demand

Before we turn to the empirics and model, we present 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-rent ratio and homeownership rate as sufficient statistics for calibration.

Figure 1: Price-Rent Ratio vs. Homeownership Rate



Notes: The figure displays national data at the quarterly frequency. The price-rent ratio is obtained from the BEA and Flow of Funds, as the ratio of the value of residential housing owned by households (FRED code: BOGZ1FL155035013Q) to the value of owner-occupied housing services (FRED code: A2013C1A027NBEA). We use an interpolation scheme to convert owner-occupied housing services from annual to quarterly (see Appendix B.5). The homeownership rate is obtained from the census (FRED code: RHORUSQ156N).

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, the relative price (price-rent ratio) the marginal renter is willing to pay to own a home, and supply, the relative price at which the marginal landlord is willing to sell that home. These equilibria were fairly stable in the pre-boom era (1965-1997), with most observations clustered in the lower left portion of the figure. However, this pattern changed dramatically during the 1997-2006 housing boom, during which 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 ending close to the typical values from the pre-boom era.

To understand what forces were behind these patterns, we present a simple supply and demand treatment that illustrates the key mechanism at work in the equilibrium model we develop in Section 5. As in Figure 1, we use the price-rent ratio on the y-axis and the homeownership rate on the x-axis. These axes thus represent the relative price and relative quantity of owned vs. rented housing. We use relative rather than absolute prices and quantities of housing to ensure that changes are driven by the rent versus own

margin rather than the construction margin.⁵

Demand for owner-occupied housing comes from constrained households who require mortgages to own. As the price-rent ratio rises, more of these households prefer renting to owning, creating a downward slope. Supply comes from landlords who can sell units of rental housing to households as owner-occupied housing. The slope of the tenure supply curve reflects the willingness of landlords to sell more units as the price-rent ratio rises, while shifts in the tenure supply curve reflect shocks that change in landlords' fundamental values of houses relative to rents.⁶ We note that this margin is distinct from changes in the absolute quantity of housing units via the construction sector.

Our supply-and-demand framework is displayed graphically in Figure 2. To begin, Panel (a) 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), Justiniano, Primiceri, and Tambalotti (2015), and Greenwald (2018), in which households cannot rent housing, corresponding to a fixed homeownership rate of 100%. In our framework, this corresponds to a perfectly inelastic tenure supply, indicated by the vertical line in Panel (a). This curve intersects the downward sloping demand curve to generate an equilibrium in price-rent versus homeownership rate space.

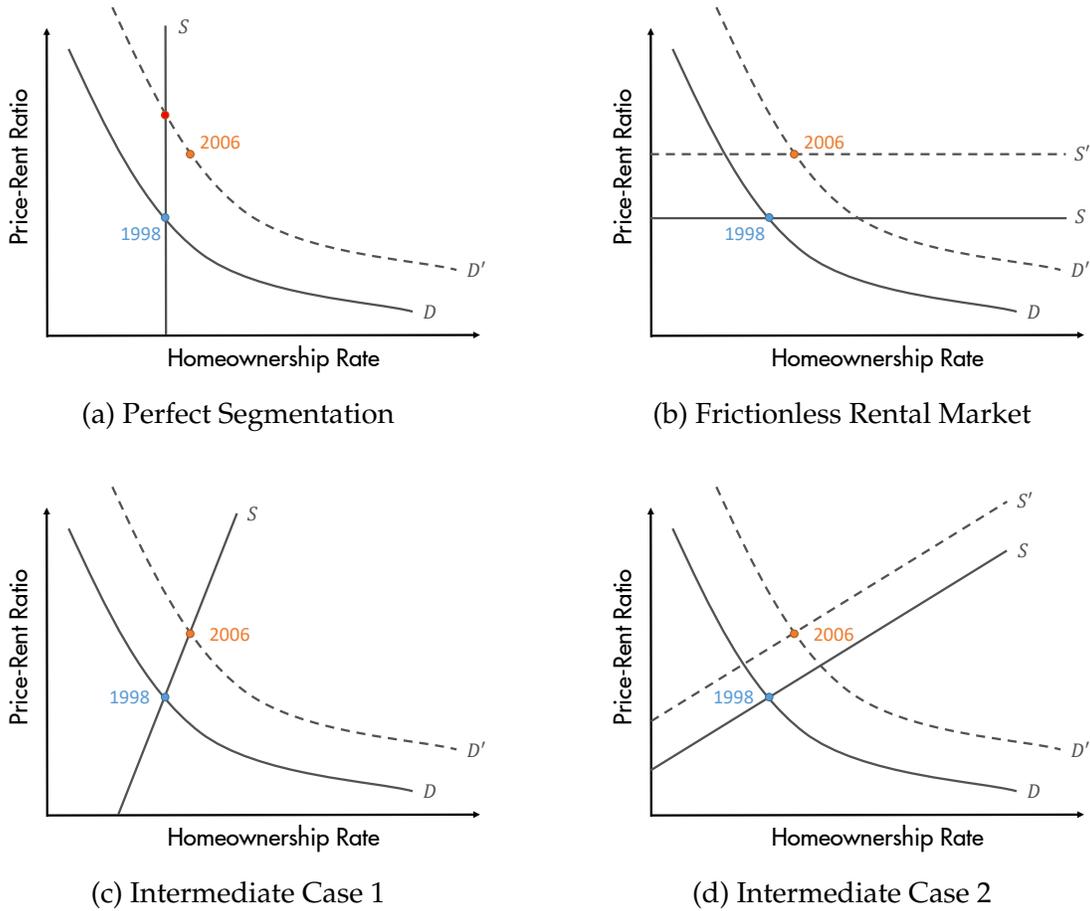
From this starting point, we can consider the impact of a credit expansion that shifts demand outward from the solid curve D to the dashed curve D' , as improvements in access to or cost of financing makes more households willing to purchase instead of renting at a given price. Despite this increase in demand, these households have no one to trade with except each other. As a result, these models imply that a relaxation of credit produced a large rise in house prices during the boom. By construction, however, they cannot reproduce the rise in homeownership displayed in Figure 1.

Panel (b) considers the alternative extreme case of 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. (2020). This specification leads to a perfectly elastic (horizontal) tenure supply curve, as landlords are willing to buy or sell an unlimited amount of housing at a price equal to the present value of rents. Since this present value does not directly depend on credit, a credit-driven expansion of demand increases the homeownership rate but not the price-rent ratio. Instead, reproducing the joint empirical

⁵For example, new housing construction typically increases the quantity of housing and decreases its price, but has no clear impact on either the price-rent ratio or the share of housing that is owner occupied. Our definitions eliminate these fluctuations which are not relevant for our analysis.

⁶If landlords require credit, a credit supply shock would also shift the tenure supply curve upward. We abstract from landlord credit in our baseline model but return to it in Section 8.

Figure 2: Supply and Demand Treatment of Model Intuitions



pattern requires a separate upward shift in the tenure supply curve, indicated by the horizontal dashed line in Panel (b). Since prices are equal to the present value of rents in this model, a shock to future rents can move prices relative to current rents, such as the shock to future rental beliefs used in Kaplan et al. (2020).⁷

While the literature to date has centered on these polar cases of perfect segmentation and zero segmentation, this paper introduces a framework that allows for intermediate levels of frictions, which corresponds to an upward-sloping tenure supply curve as in Panel (c). In this case, a credit expansion that shifts housing demand causes an increase in both the price-rent ratio and the homeownership rate as the equilibrium moves up the tenure supply curve. Panel (c) shows that such a model can in principle explain the joint empirical pattern observed during the housing boom with only a single shock. However,

⁷A shift in landlord discount rates for the same set of rental cash flows would have a similar effect.

we could also combine a flatter tenure supply curve with a shift in supply to obtain the same equilibrium movement in house prices and homeownership, as in Figure (d). Indeed, any slope of the tenure supply curve could be combined with the appropriate shift in tenure supply to generate the observed dynamics. Ultimately, to disentangle these competing explanations, we need to discipline the slope of the tenure supply curve.

We do so empirically. As is typical in the simultaneous equations literature, the slope of the tenure supply curve can be uncovered using a shock to demand. In Section 3 we use a set of credit supply shocks from the literature that provide exactly this type of variation by increasing demand for owner-occupied housing to estimate the elasticity of tenure supply. With this slope in hand, we then write down a structural model that we calibrate to this estimated slope and use the calibrated model to decompose the role of credit in the 2000s housing boom and bust.

3 Empirical Approach

Motivated by the intuition in Section 2, our goal is to estimate the slope of the tenure supply curve, equal to the ratio of the elasticity of the price-rent ratio to an identified demand shock (an expansion of credit) to the elasticity of the homeownership rate to that same shock. Doing so using ordinary least squares (OLS) is problematic because credit is endogenous and housing market conditions can naturally affect credit supply. Furthermore, credit measures may be subject to measurement error. To address these issues, we seek an instrument for credit supply.

We use three different off-the-shelf identification approaches from the literature to instrument for credit supply. While all three approaches are limited in their statistical power, particularly for homeownership rates, all three provide consistent results and thus reinforce one another. In the remainder of this section, we present our data, our basic regression framework, and then describe each empirical approach. Details on our instruments' first stages and robustness checks can be found in Appendix B.

3.1 Data

We construct an annual panel at the core-based statistical area (CBSA) level that merges together data on house prices, rents, homeownership rates, credit, and controls. The data set is slightly different for each empirical approach we use, so we describe the common data sources first, then present these variations in Sections 3.2 - 3.4. Further details on our data construction can be found in Appendix B.

For house prices, we use the CoreLogic repeat sales house price index collapsed to an annual frequency, and check robustness to using the FHFA indices in the Appendix. For credit we use Home Mortgage Disclosure Act (HMDA) data, which we collapse to the CBSA level. Our main measure of credit is the dollar volume of loan originations; in the Appendix we assess robustness using the number of loans and the loan-to-income ratio.

For rents, we use the CBRE Economic Advisors Torto-Wheaton Same-Store rent index (TW index), a high-quality repeat rent index for multi-unit apartment buildings.⁸ It is available quarterly for 53 CBSAs beginning in 1989 and 62 CBSAs beginning in 1994. Although using the TW index limits our sample sizes, it improves on rent measures typically used in the literature in two ways. First, its repeat sales methodology is preferable to median rent measures typically used in the literature, which are biased by changes in the composition of leased units. Second, while median rents tend to be sticky and slow moving due to contractual rigidities, the TW index uses asking rents on newly-leased apartments, which better reflect current market conditions.

Our homeownership data come from the Census Housing and Vacancy Survey (HVS), which provides annual estimates of the homeownership rate at the CBSA level from 1986 to 2017. These data are only available for an unbalanced panel of 82 CBSAs and are somewhat noisy as they are obtained from a supplement to the Current Population Survey with only 72,000 households nationwide. The HVS is also complicated by decennial changes in CBSA definitions. In the baseline results, we treat CBSAs where changing definitions increase or decrease the homeownership rate significantly as separate CBSAs, but present robustness checks in the Appendix dropping any CBSA that experiences a changing definition.⁹ For robustness, we supplement our data with alternative homeownership rates from the American Community Survey (ACS), which are available for a larger pool of cities than our baseline measure but begin only in 2005.

⁸CBRE EA uses data on effective rents, which are asking rents for newly-rented units net of other leasing incentives. CBRE builds a historical rent series for each building and computes 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.

⁹Specifically, we use county-level data on homeownership rates for the full population of households from the adjacent decennial Censuses. If the homeownership rate changes by more than 5%, we treat the CBSA before and after the change as different CBSAs, creating an unbalanced panel. This approach uses the full data but ensures that fixed effects and impulse responses are not affected by a jump in the homeownership rate due to CBSA definition changes.

3.2 First Empirical Approach: Loutskina and Strahan (2015)

Our first and most statistically powerful empirical approach follows LS in using a shift-share instrument based on the conforming loan limit (CLL). The CLL represents the maximum loan size eligible for securitization by Fannie Mae and Freddie Mac. Because mortgages backed by Fannie Mae and Freddie Mac receive subsidized interest rates, an increase in the CLL represents an increase in the incidence of this subsidy and a positive shock to the supply of mortgage credit for borrowers newly able to take advantage of it.

The operating principle of the instrument is that the same nationwide change in the CLL should have stronger effects in cities where a larger fraction of loans are close to this threshold since more new loans should shift from being unsubsidized to subsidized.¹⁰ For a concrete example, an average of 7.2% of loans originated in San Francisco over our sample fall within 5% of the next year’s conforming loan limit, compared to an average of only 0.4% in El Paso. Our instrument exploits the fact that a change in the CLL should therefore have a bigger average effect in San Francisco than in El Paso.

To construct an instrument that exploits this CLL variation, we follow LS and define:

$$Z_{i,t} = \left[\begin{array}{c} \text{ShareNearCLL}_{i,t} \times \%ChangeInCLL_t \\ \text{ShareNearCLL}_{i,t} \times \%ChangeInCLL_t \times \text{SaizElasticity}_i \end{array} \right].$$

To measure the share of homes with loans near CLL, we follow LS in 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. We also follow LS in using both the standard share-shift and a triple interaction with the housing supply elasticity as estimated by Saiz (2010) as instruments to allow for potential heterogeneity in the effect of the CLL change by housing supply elasticity. Since the CLL has occasionally varied by region, we use only changes in the national CLL to construct our instrument.¹¹

We then estimate the impulse response of a change in credit $C_{i,t}$ on an outcome variable of interest $Y_{i,t}$ using a local projection instrumental variables (LP-IV) approach.¹² 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).

¹⁰See Adelino et al. (2020) for an implementation based directly on house prices rather than loan sizes.

¹¹As part of the HERA legislation in 2008, Congress created more transparent procedures for changing the national CLL, and allowed the CLL to rise more in high-cost cities if their local house price index grew sufficiently quickly. This would violate our IV exclusion restriction because the change in the CLL would be mechanically correlated with lagged local outcomes. Consequently, in constructing the instrument we use the change in the national CLL regardless of the change in the local CLL in high-cost areas.

¹²In our empirical application, $C_{i,t}$ is the volume of new purchase loan origination in location i at time t ; we consider other measures of credit in the Appendix.

We extend this to the panel context and add CBSA and time fixed effects following Chen (2019). Specifically, we use two-stage least squares to estimate:

$$\log(Y_{i,t+k}) = \xi_i + \psi_t + \beta_k \Delta \log(C_{i,t}) + \theta X_{i,t} + \epsilon_{i,t} \quad (1)$$

$$\Delta \log(C_{i,t}) = \phi_i + \chi_t + \gamma Z_{i,t} + \omega X_{i,t} + e_{i,t}, \quad (2)$$

for $k = 0, \dots, 5$, where β_k is our coefficient of interest, ξ_i and ϕ_i are location fixed effects, ψ_t and χ_t are time fixed effects, and $X_{i,t}$ are controls. Standard errors are clustered by CBSA.

The formal 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. To address the potential failure of this condition due to serial correlation, we follow Ramey (2016) and Ramey and Zubairy (2018) in including two lags each of our instruments $Z_{i,t}$, our outcome variable $\log(Y_{i,t})$, and our first stage variable $\Delta \log(C_{i,t})$ as controls. We supplement these with a number of additional control variables to ensure that our estimates are based purely on the share-shift variation in our instrument. We include both CBSA effects, which absorb any average differences across areas including their supply elasticity, and time fixed effects, which absorb aggregate conditions including any average effects of the CLL on the national housing and mortgage markets. We also directly control for $ShareNearCLL_{i,t}$ and its lag so that time variation in this share is not driving our estimates.

The identifying assumption for this first empirical approach is that conditional on our controls there is no unobservable variable that varies with both the fraction of loans originated in the previous year close to the CLL (or its interaction with the Saiz elasticity) and that also varies with changes in the national CLL. For example, one might be concerned that cities with higher prices tend to be more exposed to national business cycles, and that CLL changes are also correlated with these cycles. To address such concerns, we conduct robustness checks in the Appendix demonstrating that time-varying city characteristics are not driving our results.

Since we are also interested in estimating the slope of the tenure supply curve directly, we also modify the LP-IV approach to directly estimate the impulse response of the slope. To do so, we use the homeownership rate as $Y_{i,t+k}$ and replace $\Delta \log(C_{i,t})$ with the price-rent ratio $\log(PRR_{i,t+k})$ in the first stage.¹³ The coefficients β_k then represent the inverse of the slope of the supply curve. We estimate this inverse slope rather than the slope itself

¹³Because we want to obtain the ratio of the price-rent and homeownership rate response at the same horizon, we use the price-rent ratio k periods ahead in our first stage, instead of the contemporaneous credit growth $\Delta \log(C_{i,t})$ in our previous regressions. In all cases we use the time t instrument.

because in practice the instrument has a far stronger effect on price-rent ratios than home-ownership rates, so estimating the inverse slope sidesteps weak instrument problems.

For our sample, we use an unbalanced panel of 62 CBSAs with both rents and home-ownership rates from 1992 to 2016 for our analysis, although we note that most of the identifying variation is obtained over the period 1992 to 2006 since the national CLL is fixed through the bust and rebound.¹⁴

3.3 Second Empirical Approach: Di Maggio and Kermani (2017)

Our second approach follows DK by exploiting an expansion of credit that occurred due to the OCC's 2004 preemption of state-level anti-predatory-lending laws (APLs) for national banks. States implemented APLs in 1999 to limit the terms of mortgages made to riskier borrowers. The preemption thus allowed national banks to expand credit more readily to riskier borrowers, providing a shock to the supply of credit. DK identify credit supply shocks by comparing counties with different exposures to national banks that were regulated by the OCC before and after the change. We adapt their approach to CBSAs to make use of homeownership rate data only available at this level.

We define the DK instrument as:

$$Z_i = APL_{2004} \times OCC_{2003},$$

where APL_{2004} is an indicator for whether the state that the majority of the CBSA resides in has an anti-predatory-lending law by 2004, and OCC_{2003} is the share of mortgage originated by OCC-regulated banks in 2003, obtained from HMDA data.

Because the instrument only induces variation across cities in response to a single credit supply event, we cannot use the LP-IV approach described above, which would require variation in the instrument both across cities and over time. Instead, we follow DK and in using an event study approach. We also follow much of the literature in focusing on the reduced form, regressing the outcome variables directly on the instrument. This is sufficient to obtain the slope of the supply curve, but implies that we cannot interpret coefficient magnitudes in units of credit.

In particular, we estimate the regression:

$$\log(Y_{i,t}) = \xi_i + \psi_t + \sum_{k \neq \tau} \beta_k Z_i \mathbf{1}_{t=k} + \theta X_{i,t} + \epsilon_{i,t}. \quad (3)$$

¹⁴In practice, the CLL never adjusts downward, so it typically remains flat during housing downturns until prices pass their previous peak.

The coefficients β_k represent the reduced form effect of the instrument at each date in time relative to a base period τ for which β is normalized to zero. To ensure that only the interaction of APL_{2004} and OCC_{2003} is used for identification, we follow DK in controlling for both of these variables directly in addition to including CBSA and year fixed effects (ξ_i and ψ_t). Our controls ($X_{i,t}$) include the lag of the outcome variable as well as all additional controls used by DK in their original analysis, with the exception of a proprietary measure of the share of loans originated to subprime borrowers that is not crucial for their results. Since this identification strategy is similar to a difference-in-difference with Z_i measuring exposure, a key test of the identifying assumptions is that there are parallel pre-trends or equivalently that β is equal to zero prior to date τ .

The identifying assumptions are similar to a differences-in-differences approach: There must be parallel trends in the absence of treatment. DK provide extensive support for this identifying assumption in their paper, and we replicate this finding of no pre-trends prior to 2003 in our empirical specifications.

We follow DK in estimating equation (3) using growth rates for $Y_{i,t}$, so that the outcome variable is the log change in house prices or homeownership rates. To obtain an impulse response in levels, we then add up the coefficients of interest from the base period to each indicated period and compute standard errors by the delta method.

DK kindly provided us with their data set, and we use their data directly collapsed to the CBSA level to be as consistent with their paper as possible. We then merge in CBSA-level CoreLogic house price index and census homeownership rates. This yields 262 CBSAs from 2001 to 2010 for house prices and 82 CBSAs from 2001 to 2010 for homeownership rates. Due to the limited power of a single event study, we focus on house prices rather than price-rent ratios, which would require cutting our sample further to the subset with available rent data. We follow DK by weighting our regressions by population and by clustering standard errors by CBSA. For homeownership rates, we drop cities that have a change in Census homeownership definitions in 2005.

3.4 Third Empirical Approach: Mian and Sufi (2019)

Our third approach exploits 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 MS and Nadauld and Sherlund (2013). MS 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. MS argue that the PLS expansion had a larger effect on lending by mortgage lenders that rely on non-core de-

posits to finance mortgages, measured at the bank level as the ratio of non-core liabilities to total liabilities (NCL). They hypothesize that NCL banks, which are funded less by deposits, should have a greater exposure to the PLS expansion, and show that high NCL banks did in fact expand their lending more after following roughly parallel trends prior to 2002. Like the DK instrument, this approach isolates a non-prime credit supply shock, in contrast to the LS shock that directly affects prime borrowers only.

The MS instrument is defined as:

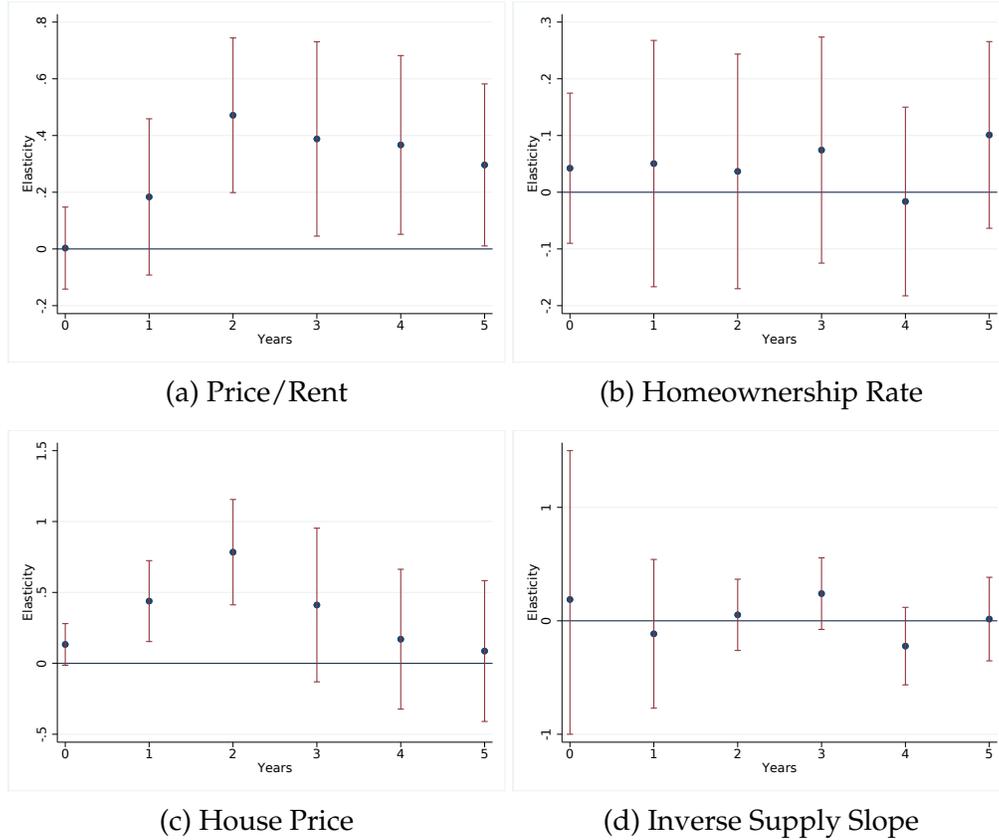
$$Z_i = NCLShare_i^{2002},$$

where $NCLShare_i^{2002}$ is MS's measure of CBSA-level exposure to high NCL lenders, equal to the origination-weighted average of lender-level NCLs in a CBSA based on 2002 originations. MS argue that the city-level NCL exposure satisfies the relevant exclusion restriction and is a valid instrument for credit. Because this instrument also induces variation across cities in response to a single event, we use the same reduced form event study approach (3) that we use for the DK instrument. We follow MS in weighting by the number of housing units and including year and CBSA fixed effects, and cluster at the CBSA level.

The MS instrument is underpowered using only the CBSAs for which we have Census HVS homeownership rates. Consequently, for this empirical approach we expand our data set by using ACS data for homeownership rates and FHFA data for house prices. This ACS-FHFA data sample covers 245 CBSAs from 1990 to 2017 for prices and 245 CBSAs from 2005 to 2017 for the homeownership rate. However, this means that we must use house prices in place of our preferred outcome variable, the price-rent ratio.¹⁵ Using this data sample also prevents us from setting the base year of 2002 used by MS because our ACS homeownership rate data begins in 2005. We instead use 2013 as the base year, since our estimates imply that the house price response returned to its 2002 level in 2013. Our results are robust to using peak-to-trough changes rather than a particular base year.

¹⁵We do not use the ACS homeownership rates for the LS 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 rather than new rents. Due to the long term nature of leases, average rents move much less than new rents, so a price-rent ratio constructed with ACS data looks nearly identical to the same regression using prices as an outcome.

Figure 3: Loutskina-Strahan Instrument LP-IV Impulse Responses



Notes: 95 % confidence interval shown in red bars. The figure shows panel local projection instrumental variables estimates of the response of the indicated outcomes to dollar credit volume. For panels (a) to (c), the first and second stages are indicated in equations (2) and (1), respectively. The two instruments are $ShareNearCLL_{i,t} \times \%ChangeInCLL_t$ and $ShareNearCLL_{i,t} \times \%ChangeInCLL_t \times Z(SaizElasticity_i)$. Control variables include $ShareNearCLL_{i,t}$ and its lag, lags of both instruments, and two lags of both the outcome variable and the endogenous variable. For panel (d), use the homeownership rate as our outcome variable and replace the log credit growth at time t with the log price-rent ratio at time $t + k$ to obtain a coefficient for the inverse supply curve slope. All standard errors are clustered by CBSA.

4 Empirical Results

4.1 Loutskina-Strahan LP-IV Results

Figure 3 plots the impulse responses of our outcome variables to credit growth (β_k in equation (1)) using LP-IV and the LS instrument. Beginning with Panel (a), we observe that the price-rent ratio rises for two years following the shock before peaking at an increase of 0.471. Our estimates are significant at the 5% level in years 2 through 5. The smaller and statistically insignificant responses in years 0 and 1 are typical of house price dynamics, which exhibit sluggish reactions and short run momentum (Guren (2018)).

Decomposing this result, this behavior of the price-rent ratio is due to a larger hump-shaped response of house prices, which peak at 0.784 after two years (Panel (c)), and a hump-shaped response of rents, which reaches 0.160 at the same horizon (Appendix Figure B.1). Our results for the effect of credit on house prices are consistent with those found in the literature, such as Glaeser et al. (2012), Adelino et al. (2020), Favara and Imbs (2015), and Di Maggio and Kermani (2017). For instance, Favara and Imbs (2015) find an elasticity of house prices to loan volumes over one year of 0.134, which is extremely close to our estimate of 0.133 at the same horizon.¹⁶

In contrast to our results on price-rent ratios and house prices, we find no significant response of homeownership rates to credit shocks. While this is in part due to lower statistical power stemming from a noisier data set, the point estimates are also consistently small, reaching 0.037 after two years, and peaking at 0.101 after 5 years. For a naive “back-of-the-envelope” measure of the relative slope, we can simply divide the point estimates in Panel (a) by those in Panel (b) to obtain ratios of 12.83 at the 2-year horizon, 5.22 at the 3-year horizon, -22.50 at the 4-year horizon, and 2.93 at the 5-year horizon, corresponding to a range of slopes between 2.9 and infinity.¹⁷

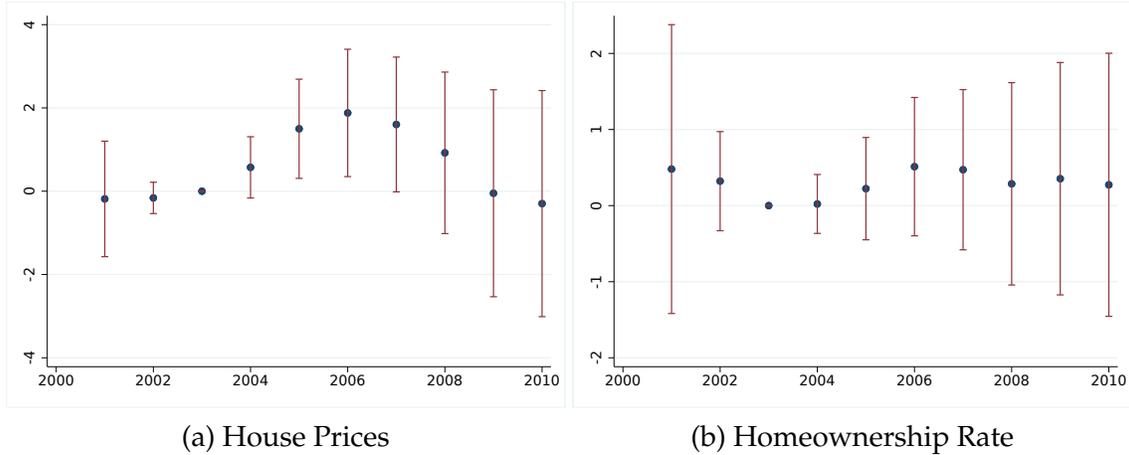
Beyond these naive ratios, we pursue a more econometrically precise approach by directly estimating the inverse of this ratio. We reestimate our IV specification (1) - (2) using the homeownership rate as the outcome variable $Y_{i,t+k}$ and the price-rent ratio in period $t + k$ in place of log credit growth as the independent variable. As described in Section 3, we estimate the *inverse* slope (response of homeownership relative to response of price-rent) because our imprecise results in Figure 3 Panel (b) would yield a weak first stage if the homeownership rate were used as the independent variable. Panel (d) shows that the inverse slope is small and not statistically different from zero, with point estimates of 0.05 at the 2-year horizon, 0.24 at the 3-year horizon, -0.22 at the 4-year horizon, and 0.02 at the 5-year horizon.¹⁸ Inverting these estimates yields supply curve slopes between between 4.2 and infinity. The upper bounds of the 95% confidence intervals for the inverse slope are 0.37 at the 2-year horizon, 0.56 at the 3-year horizon, 0.12 at the 4-year horizon, and 0.38 at the 5-year horizon, corresponding to lower bound estimates of the (non-inverted) slope between 1.8 and 8.4.

¹⁶We find a larger response of rents than Favara and Imbs, likely because we are using the TW rent index, which provides the rent of a newly-rented multi-family unit using a repeat sales methodology, rather than stickier median rents as used by prior literature.

¹⁷Since a downward sloping supply curve is implausible, negative inverse ratios are best interpreted as infinite (perfectly inelastic) slopes, since these offer the smallest possible ratio of the homeownership rate response to the price-rent ratio response.

¹⁸We omit the ratios for the 0-year and 1-year horizons since Panel (a) implies that our first stage is not statistically significant at these horizons.

Figure 4: Di Maggio-Kermani APL Preemption Reduced Form



Notes: 95 % confidence interval shown in red bars. Each panel shows estimates of the cumulative sum from 2003 of β_k for each indicated year estimated from equation (3), with the instrument being $Z_i = APL_{2004} \times OCC_{2003}$ and 2003 being the base year. The regression is weighted by population and standard errors are clustered by CBSA. The controls include median income growth, population growth, the Saiz (2010) elasticity interacted with a dummy for post-2004, the fraction of loans originated by HUD-regulated lenders interacted with a dummy for post-2004, and the fraction of HUD-regulated lenders interacted with a dummy for APLs. The controls are as Di Maggio and Kermani (2017) except our data is at the CBSA level and omits a control for the fraction of loans originated that are subprime (FICO under 620), which is based on proprietary data. All regressions are weighted by population and standard errors are clustered by CBSA, as in the original DK paper.

4.2 Di Maggio-Keramani APL Preemption Results

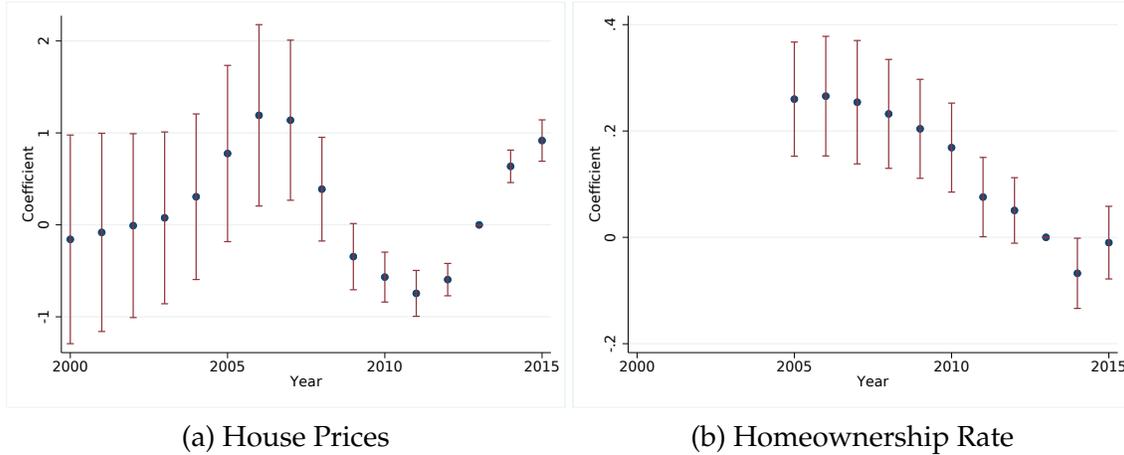
Figure 4 shows the results for the reduced form event study using the DK instrument. As a reminder, we estimate equation (3) with house prices as the outcome in log changes and then cumulate the coefficient β coefficients from the base period to the indicated period to obtain an IRF in levels.

Panel (a) displays the response of house prices for the sample of CBSAs for which we have homeownership rate data. This differs slightly from the original DK estimation, which used a broader sample of cities.¹⁹ The impulse response shows no significant pre-trends are evident prior to 2003. After 2004, the results demonstrate a classic hump-shaped impulse response for house prices peaking in 2007 at 1.60, which is significant at the 5% level. Panel (b) of Figure 4 shows the impulse response for homeownership rates is generally smaller, peaking in 2006 at 0.51, and is far from statistically significant. A simple division of the values yields slopes of 6.72 in 2005, 3.67 in 2006, and 3.40 in 2007.²⁰

¹⁹Using the broader sample of CBSAs we are able to largely reproduce DK's results (see Figure B.10). Cities with homeownership rate data tend to be larger, more inelastic cities, leading us to find larger responses of house prices than in the original DK estimation.

²⁰As in Section 3.2, we focus on slopes in periods with house prices responses following the process of

Figure 5: Mian-Sufi PLS Expansion Reduced Form



Notes: 95 % confidence interval shown in red bars. Panels A and B shows estimates of the effect of a city's NCL share on each outcome based on estimating equation (3) with the instrument being $Z_i = NCLShare_i^{2002}$ and 2013 being the base year. For panel C, we use the same controls but the outcome variable is the homeownership rate and the credit variable is replaced with log house prices to obtain a coefficient for the inverse supply curve slope. All standard errors are clustered by CBSA as in the original MS paper.

4.3 Mian-Sufi PLS Expansion Results

Figure 5 displays our estimated β_k coefficients by year for the Mian and Sufi NCL share instrument. As a reminder, while the PLS market expansion is in 2002, we have normalized 2013 to be the base year due to data availability, and also use house prices as the outcome due to power concerns. Panel (a) shows a zero effect on prices prior to 2002, followed 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 peaks in 2006 at 1.19 while the effect on the homeownership rate also peaks in 2006 at 0.27. A simple division of these point estimates yields a slope of 4.49, while the same exercise using the 2007 coefficients yields a ratio of 4.48.

4.4 Summary

While all three of our empirical strategies have limitations, they each find evidence for a slope of the supply curve of at least three, and often higher. We obtain these consistent results even though our instruments rely on completely distinct sources of variation and influence different segments of the mortgage market, with the LS instrument affecting

sluggish adjustment, as these provide the best analogue to the corresponding moment in the model. We similarly report years satisfying these criteria for our MS results in Section 4.3 below.

credit to prime borrowers and the DK and MS instruments affecting credit to more risky borrowers. Our finding of broad agreement across specifications leads us to conclude that our general finding that house prices response more than homeownership rates in response to a shock to credit supply is robust. In Section 6, we use a model to provide economic interpretation of these numeric slope values, showing that our estimates provide clear lower bounds on the degree of segmentation needed to match the data.

5 Model

This section develops an equilibrium model that we use to quantitatively evaluate the role of credit in driving house prices, with a focus on the 2000s boom-bust cycle.

Demographics. There is a representative borrower, landlord, and saver, denoted B , L , and S , respectively. Each is infinitely lived. We assume perfect risk sharing within each type, allowing for aggregation to a representative agent of each type.

Housing Technology. Housing is produced by construction firms (described below) whose supply at the end of period t is denoted \bar{H}_t . Housing can be owned either by borrowers, by savers, or by landlords, where landlords 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 its stock, and is sold ex-dividend (i.e., after the service flow is consumed). It requires a per-period maintenance equal to fraction δ of its current value.

While housing is traded by borrowers and landlords, our main specification imposes that savers 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. In Section 8, however, we relax this assumption to allow savers to freely trade housing with borrowers. For notational convenience, we denote the stock of housing not owned by savers, and therefore ultimately inhabited by borrowers (as either owners or renters), as $\hat{H}_t = \bar{H}_t - \bar{H}_S$.

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

$$U_j = \sum_{t=0}^{\infty} \beta_j^t \log \left(c_{j,t}^{1-\xi} h_{j,t}^{\xi} \right), \quad j \in \{B, S\}$$

where c is nondurable consumption, and h is housing services. We assume that landlords are risk neutral and maximize:

$$U_L = \sum_{t=0}^{\infty} \beta_L^t c_{L,t}.$$

Risk neutrality nests typical specifications in the literature, which often model landlords as a foreign-owned, profit-maximizing firm (see e.g., Kaplan et al. (2020)).

Asset Technology. Borrowers and landlords can trade long-term mortgage debt with savers in 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 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 $(r^* + \nu)M^*$, $(1 - \nu)(r^* + \nu)M^*$, $(1 - \nu)^2(r^* + \nu)M^*$, etc. Mortgage loans are prepayable, with exogenous fraction ρ prepaying their loans in a given period, and are also nominal, meaning that real balances decay each period at the constant rate of inflation π .

As in Greenwald (2018), the average size of new loans for borrower i (denoted $M_{i,t}^*$) is subject to both loan-to-value (LTV) and payment-to-income (PTI) limits at origination:

$$M_{i,t}^* \leq \theta^{LTV} p_t H_{i,t}^* \quad (4)$$

$$M_{i,t}^* \leq \frac{(\theta^{PTI} - \omega) \text{income}_{i,t}}{r_{B,t}^* + \nu + \alpha}, \quad (5)$$

where p_t is the price of housing, $H_{i,t}^*$ is the borrower's new house size, and ω and α are offsets used to account for non-housing debts, and taxes and insurance, respectively.

Ownership Benefit Heterogeneity. Without additional heterogeneity, the model would be unable to generate a fractional and time-varying homeownership rate. 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 thus 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 introduce 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 $\omega_{i,t}$ times the market rent for that unit, where $\omega_{i,t}^B \sim \Gamma_{\omega,B}$ is drawn i.i.d. across borrowers and time. Symmetrically, if a landlord owns unit i of housing, he or she receives surplus equivalent to $\omega_{i,t}^L \sim \Gamma_{\omega,L}$ times the market rent for that unit. Because we perceive these benefits and costs, particularly those of the borrower, as likely non-financial, we rebate them lump-sum to households, so that they do not have any effect on the resource constraint in equilibrium. Since we apply borrower heterogeneity at the household level but landlord heterogeneity at the property level, the two dimensional sorting problem is trivial: all properties with sufficiently low $\omega_{i,t}^L$ are owned, and they are owned by the households with the largest $\omega_{i,t}^B$.

There are several forms of heterogeneity that would map intuitively into this framework. On the borrower side, heterogeneity in the value of ownership likely stands in for household age, family composition, ability to make a down payment, and true personal preference for ownership. On the landlord side, we conjecture that the biggest source of heterogeneity is on the suitability of different properties for rental, as documented for instance by Halket, Nesheim, and Oswald (2020). 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. Under this interpretation, as the homeownership rate rises, the marginal converted property is easier to convert and maintain, and is valued more highly by landlords relative to the rent it produces, implying an upward-sloping tenure supply curve.

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

Borrower's Problem. The borrower maximizes expected lifetime utility subject to the borrowing constraints (4), (5), and the budget constraint:

$$c_{B,t} \leq \underbrace{(1-\tau)y_{B,t}}_{\text{after-tax income}} + \underbrace{\rho_B \left(M_{B,t}^* - \pi^{-1}(1-\nu_B)M_{B,t-1} \right)}_{\text{net mortgage iss.}} - \underbrace{\pi^{-1}(1-\tau)X_{B,t-1}}_{\text{interest payment}} - \underbrace{\nu_B \pi^{-1}M_{B,t-1}}_{\text{principal payment}}$$

$$-\underbrace{\rho_B p_t (H_{B,t}^* - H_{B,t-1})}_{\text{net housing purchases}} - \underbrace{\delta p_t H_{B,t-1}}_{\text{maintenance}} - \underbrace{q_t (h_{B,t} - H_{B,t-1})}_{\text{rent}} + \underbrace{\left(\int_{\bar{\omega}_{B,t-1}} \omega d\Gamma_{\omega,B} \right) q_t \hat{H}_{t-1}}_{\text{owner surplus}} + \underbrace{T_{B,t}}_{\text{rebates}},$$

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, $\bar{\omega}_{B,t} = \Gamma_{\omega,B}^{-1}(1 - H_{B,t}/\hat{H}_t)$, 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 mortgage interest payments are tax deductible.

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

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

Landlord's Problem. The landlord's problem is similar to that of the borrower, with two key exceptions: (i) the landlord sells housing services to the borrower instead of consuming them, and (ii) the landlord does not use credit — an assumption we relax in Section 8. The landlord is subject to the budget constraint:

$$c_{L,t} \leq \underbrace{(1 - \tau)y_{L,t}}_{\text{after-tax income}} - \underbrace{p_t (H_{L,t} - H_{L,t-1})}_{\text{net housing purchases}} - \underbrace{\delta p_t H_{L,t-1}}_{\text{maintenance}} + \underbrace{q_t H_{L,t-1}}_{\text{rent}} + \underbrace{\left(\int_{\bar{\omega}_{L,t-1}} \omega d\Gamma_{\omega,L} \right) q_t \hat{H}_{t-1}}_{\text{owner surplus}} + \underbrace{T_{L,t}}_{\text{rebates}},$$

and the market clearing condition $\bar{\omega}_{L,t} = \Gamma_{\omega,L}^{-1}(1 - H_{L,t}/\hat{H}_t)$.

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

$$\begin{aligned} c_{S,t} &\leq \underbrace{(1 - \tau)y_{S,t}}_{\text{after-tax income}} - \underbrace{p_t (H_{S,t}^* - H_{S,t-1})}_{\text{net housing purchases}} - \underbrace{\delta p_t H_{S,t-1}}_{\text{maintenance}} + \underbrace{T_{S,t}}_{\text{rebates}} \\ &\quad + \underbrace{\pi^{-1}(\bar{r}_B + \nu_B)M_{B,t-1}}_{\text{total payment}} - \underbrace{\rho_B \left(\exp(\Delta_{B,t})M_{B,t}^* - \pi^{-1}(1 - \nu_B)M_{B,t-1} \right)}_{\text{net mortgage iss.}}, \end{aligned}$$

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 the mortgage rate exceeds the rate on a risk-free bonds with the payment structure, allowing for exogenous variation in mortgage spreads. In addition to the budget constraint, the saver must also satisfy the fixed housing demand constraint $H_{S,t} = \bar{H}_S$ at all times.

Construction Firm's Problem. We assume that new housing is produced by competitive construction firms. Following Favilukis et al. (2017) and Kaplan et al. (2020), we assume that housing is produced using nondurables Z_t and land L_t according to the technology:

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

We assume that \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 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. We note, however that these assumptions are not central to our results, as the dynamics of price-rent ratios, homeownership, and credit would be very similar in a model with a fixed housing stock (see Appendix Figure A.2).

Equilibrium. A competitive equilibrium economy consists of endogenous states $(H_{B,t-1}, M_{B,t-1}, X_{B,t-1}, \bar{H}_{t-1})$, borrower controls $(c_{B,t}, h_{B,t}, M_{B,t}^*, H_{B,t}^*)$, landlord controls $(c_{L,t}, H_{L,t})$, saver controls $(c_{S,t}, M_{S,t}^*)$, construction firm controls (L_t, Z_t) , and prices $(p_t, q_t, r_{B,t}^*)$ that jointly solve the borrower, landlord, saver, and construction firm problems, as well as the market clearing conditions:²¹

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

²¹In a slight abuse of notation we allow both the saver and borrower to choose $M_{B,t}^*$ as controls, and implicitly impose that these values must be equal in equilibrium.

5.1 Key Equilibrium Conditions

We now present the key equilibrium conditions of the model, while reserving the full set of equilibrium conditions to Appendix A.1. These key equations are the optimality conditions for borrower and landlord housing, respectively which correspond to the inverted demand and tenure supply curves:

$$p_t^{\text{Demand}}(H_{B,t}) = \frac{E_t \left\{ \Lambda_{B,t+1} \left[(1 + \bar{\omega}_{B,t}) q_{t+1} + (1 - \delta - (1 - \rho_B) \mathcal{C}_{B,t+1}) p_{t+1} \right] \right\}}{1 - \mathcal{C}_{B,t}} \quad (6)$$

$$p_t^{\text{Supply}}(H_{B,t}) = E_t \left\{ \Lambda_{L,t+1} \left[(1 + \bar{\omega}_{L,t}) q_{t+1} + (1 - \delta) p_{t+1} \right] \right\}. \quad (7)$$

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 $H_{L,t} = \hat{H}_t - H_{B,t}$ of housing.

These are standard asset pricing equations that state that the asset price is equal to the expected future payoff discounted by the relevant stochastic discount factors, here $\Lambda_{B,t+1}$ for borrowers and $\Lambda_{L,t+1}$ for lenders. The supply schedule (7) sets the current price equal to the present value of the next period cash flow (rent) for the marginal landlord $(1 + \bar{\omega}_{L,t}) q_{t+1}$ plus the next period value of the housing after maintenance costs. The demand schedule (6) is similar, but is influenced by the ability of borrower housing to collateralize debt, which is valued by borrowers. This enters through the marginal collateral value term $\mathcal{C}_{B,t}$, which represents the shadow value of the additional credit that can be collateralized by an additional dollar of housing (see Section A.1 for details).²² A relaxation of credit standards or a decrease in the cost of credit allows housing to collateralize more or cheaper credit, raising this marginal value $\mathcal{C}_{B,t}$, and increasing the reservation price. As a result, changes in credit conditions directly shift the demand schedule (6).

These equations map directly into the supply and demand framework of Section 2, where the demand and tenure supply schedules map into the $\bar{\omega}_{j,t}$ terms. Recall that these terms are defined by $\bar{\omega}_{B,t} = \Gamma_{\omega,B}(1 - H_{B,t}/\hat{H}_t)$ and $\bar{\omega}_{L,t} = \Gamma_{\omega,L}(1 - H_{L,t}/\hat{H}_t)$. As $H_{B,t}$ increases, so does the fraction of owner-occupied housing. This pushes down $\bar{\omega}_{B,t}$, as the marginal household becomes increasingly less suited for ownership, 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. In equi-

²²The exact definition is $\mathcal{C}_{B,t} \equiv \mu_{B,t} \frac{\Gamma_{\omega,B}^{LTV} \theta^{LTV}}{F_t^{LTV}}$ as in Greenwald (2018). An extra dollar of housing can collateralize θ^{LTV} of new debt for an LTV-constrained borrower, of which there are fraction F_t^{LTV} . Finally, the Lagrange multiplier $\mu_{B,t}$ on the borrowing constraint converts from the quantity of new credit to the *value* of that credit from the borrower's perspective.

librium, the level of owner-occupied housing $H_{B,t}$ adjusts so that $p_t^{\text{Demand}} = p_t^{\text{Supply}}$, and the market clears. 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, which governs the slopes of the demand and supply curves, respectively.

6 Model Quantification

We calibrate our model at quarterly frequency, with the full set of parameters displayed in Table 1. We calibrate all parameters except for $\sigma_{\omega,L}$ in Section 6.1, then calibrate $\sigma_{\omega,L}$ to directly match our empirical regressions in Section 6.2.

6.1 Main Calibration

Demographics and Preferences To determine the borrower population share, we use the 1998 Survey of Consumer Finances. In the model, borrowers are constrained households whose choice between renting and owning is influenced by credit conditions. Correspondingly, we identify a household as a “borrower” in the data if it either (i) owns a home and its 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 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 $\chi_L \rightarrow 0$ and assume that landlords do not receive labor income, instead subsisting entirely on their rental earnings.²³

For preferences, the key parameter is the borrower’s discount factor, β_B , which determines the level latent demand for credit in the economy, and in turn, how much a relaxation of credit will influence household demand for owned housing. 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.²⁴ We choose β_B so that the typical borrower would be indifferent between receiving a loan at 80% LTV, and paying the exact FHA insurance scheme for a loan at 95% LTV: an up front fee of

²³Because landlord utility is linear in consumption, assumptions about their income and consumption have essentially no impact on the results.

²⁴For example, 37.7% of Fannie Mae purchase loans required PMI over the 1999-2008 boom period (source Fannie Mae Single Family Dataset).

Table 1: Parameter Values: Baseline Calibration (Quarterly)

Parameter	Name	Value	Internal	Target/Source
<i>Demographics and Preferences</i>				
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	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	0.2	N	Davis and Ortalo-Magné (2011)
Saver housing demand	\bar{H}_S	5.299	Y	Steady state optimum
<i>Ownership Benefit Heterogeneity</i>				
Landlord het. (location)	$\mu_{\omega,L}$	-0.109	Y	Avg. homeownership rate
Landlord het. (scale)	$\sigma_{\omega,L}$	2.877	Y	Empirical elasticities (Section 6)
Borr. het. (location)	$\mu_{\omega,B}$	0.217	Y	Borr. VTI (1998 SCF)
Borr. het. (scale)	$\sigma_{\omega,B}$	0.319	Y	Implied subsidy (see text)
<i>Technology and Government</i>				
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
<i>Mortgage Contracts</i>				
Refinancing rate	$\bar{\rho}$	0.034	N	Greenwald (2018)
Loan amortization	ν	0.45%	N	Greenwald (2018)
Borr. LTV Limit	θ_B^{LTV}	0.85	N	Greenwald (2018)
Borr. PTI Limit	θ_B^{PTI}	0.36	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

1.75% of the loan, plus a spread of 80 basis points.²⁵

For the other preference parameters, we assume a standard consumption weight parameter of $\xi = 0.2$ 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 Bonds 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,

²⁵We choose the FHA scheme as it is much simpler to implement in the model than the GSE scheme, where pricing is less transparent, and insurance premia are only paid until the borrower's LTV drops below 80%. Goodman and Kaul (2017) show that the overall costs of the two forms of insurance are similar.

it is at the correct “long run” equilibrium value. Last, we set the landlord discount rate β_L to be equal to β_B . This calibration ensures that both borrowers and landlords discount future housing services and rental cash flows at essentially equal rates. As a result, shocks that shift the path of future rents will affect borrowers and landlords symmetrically, and have little impact on the equilibrium homeownership rate.²⁶

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 with c.d.f.

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

The landlord dispersion parameter $\sigma_{\omega,L}$ is the key parameter in our model, and is calibrated to directly match our empirical estimates from Section 4. As this procedure is somewhat involved, we defer further comment to Section 6.2, which is entirely devoted to the calibration of $\sigma_{\omega,L}$.

Turning to the other parameters, we set the average level of landlord ownership utility ($\mu_{\omega,L}$) to attain the correct homeownership rate among “borrowers” in the 1998 SCF (49.64%). Since all savers own in the model, this ensures an overall homeownership rate of 68.50% at steady state, matching the 1998 SCF. We next set the average level of borrower ownership utility ($\mu_{\omega,B}$) to target the average ratio of home value to income among borrowers who own homes in the 1998 SCF, equal to 8.81 (quarterly).

Last, we calibrate borrower ownership dispersion $\sigma_{\omega,B}$, which determines the rate at which borrower households switch between owning and renting as the price of housing changes. Because $\sigma_{\omega,B}$ must be identified from shifts in tenure supply (shocks to house prices exogenous to borrower demand), our empirical estimates in Section 4 using shifts in demand are not informative about this parameter. Instead, we use the empirical estimates of Berger, Turner, and Zwick (2020), who study the impact of the First Time Home Buyer credit, a 2009 - 2010 policy that subsidized the purchase of housing by up to 10% (capped at \$8,000) for renter households that did not own in the three prior years.

Berger et al. (2020) estimate that this subsidy led 3.2% of eligible renters to switch to ownership during the policy window (February 17, 2009 to July 1, 2010), a rate of 0.64% per quarter. For our calibration, we choose $\sigma_{\omega,B} = 0.319$ so that exactly 0.64% of renters

²⁶For example, if borrower and landlord discount factors differ, shocks to expectations about future rents will cause a large shift in ownership toward the type with a higher discount factor. In the absence of further disciplining information, we seek a calibration that avoids these shifts.

(equal to $0.64\% / 3.4\% = 18.8\%$ of “active” renters) would switch from renting to owning if their housing purchases were given a 10% subsidy. While this calibration strategy ignores some of the more nuanced details around this episode, the borrower dispersion parameter $\sigma_{\omega,B}$ is much less important for our results than the landlord dispersion parameter $\sigma_{\omega,L}$.²⁷ This is because our core experiments shift the demand curve and then travel along the supply curve, making the supply slope much more influential than the demand slope (see Appendix Figure A.1 for sensitivity analysis).

Technology and Government. For the construction technology, we set the amount of new land permits issued per period so that residential investment Z_t is 5% of total output in steady state. For the land weight in the construction function φ , we note that $\varphi / (1 - \varphi)$ 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.

6.2 Calibration of Landlord Heterogeneity to Our Empirical Results

We calibrate the dispersion in the landlord’s ownership cost $\sigma_{\omega,L}$ — the model’s key parameter governing the slope of credit supply and the response of house prices to a change in credit conditions — so that the model is able to reproduce as closely as possible the responses to an identified credit shock that we estimate in Section 4. In particular, we target our estimates from the LS IRF of the price-rent ratio and homeownership rate displayed in Panels (a) and (b) of Figure 3. We focus on the LS results because they are the most statistically precise, stem from a shock that is more straightforward to map into the model, and produce responses that are measured using the ideal price-rent ratio outcome variable, rather than the house price outcomes used with the DK and MS results.

Since a change in the CLL effectively adds a subsidy to the newly included mortgages, the model analogue of the empirical regression is a linearized impulse response following a shock to the mortgage spread $\Delta_{B,t}$, which we assume follows an AR(1). Once we

²⁷Such details include that not all households received the full subsidy due to income and value caps, the finding by Berger et al. (2020) that the policy increased house prices, offsetting part of the subsidy, and the potential that the share of “active” households in this episode differed from our steady state level of 3.4%.

compute our model impulse responses at quarterly frequency, we take annual averages to map our model impulse responses into the annual frequency used in our empirical regression. We then choose $\sigma_{\omega,L}$ along with the size and persistence of the shock to the mortgage spreads $\Delta_{B,t}$ corresponding to the change in the conforming loan limit to minimize the distance between the model and data impulse responses. The size and persistence of the shock to mortgage spreads can be considered to be nuisance parameters that are not of direct interest for our experiments but that are important to pin down the scale and temporal shape of the shock response, which are not directly identified by $\sigma_{\omega,L}$.

We set these three parameters to minimize the squared error between the model and the data scaled by the statistical uncertainty around the empirical estimates:

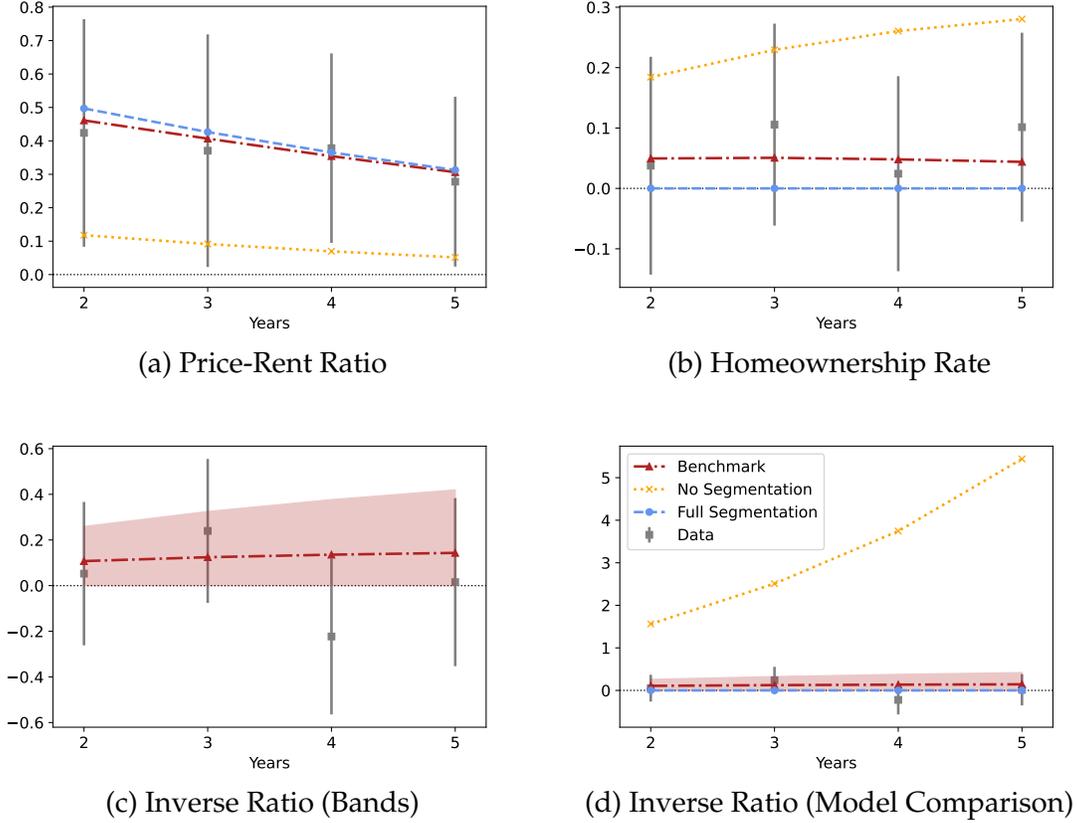
$$Q = \sum_{v \in \{PR, HOR\}} \sum_{k=2}^5 \left(\frac{IRF_{v,k}^{Model} - \beta_{v,k}}{SE_{v,k}} \right)^2, \quad (8)$$

where $IRF_{v,k}^{Model}$ is the model impulse response for variable v (either the price-rent ratio or the homeownership rate), $\beta_{v,k}$ is the corresponding estimate from Figure 3, and $SE_{v,k}$ is the estimated standard error of $\beta_{v,k}$. We restrict our estimation to the response at horizons from two years to five years. We do so because our model produces responses that jump on impact. This is typical in models of frictionless housing markets but contrasts with the data, where house prices and price-rent ratios typically display hump-shaped responses and momentum due to search and other frictions. Including the first two periods in our estimation would thus lead the model to “compromise” by systematically understating price-rent ratio growth at horizons of two to five years, in an attempt to reduce errors the first two periods. To avoid this bias due to misspecification, we therefore only ask the model to fit the data in periods after the two-year peak.

Our procedure estimates a landlord cost dispersion ($\sigma_{\omega,L}$) of 2.877, a mortgage spread shock persistence of 0.965, and a mortgage spread shock size of -0.041, where a negative shock size captures that spreads fall due to the subsidy. These estimates indicate a persistent but non-permanent shock and an annualized CLL subsidy of 17bp, which falls in the typical range of 10bp - 24bp found in the literature (Adelino et al., 2020).

To interpret our estimate of $\sigma_{\omega,L}$, Figure 6 displays our estimated empirical IRFs alongside the IRFs obtained from a model at our minimum-distance estimate (henceforth the “Benchmark” model), as well as two polar economies: one with “Full Segmentation”, corresponding to $\sigma_{\omega,L} \rightarrow \infty$, and one with “No Segmentation,” corresponding to $\sigma_{\omega,L} = 0$. These polar economies correspond to the perfectly inelastic and perfectly elastic tenure supply examples in Section 2, respectively. To isolate the role of our main parameter, we

Figure 6: Impulse Responses: Model vs. Data



Notes: Grey squares and bands indicate our benchmark LS estimates from Figure 3 and their 95% confidence intervals. These empirical estimates are plotted alongside the corresponding outputs of each model. For each model, we compute a quarterly impulse response, then average over each year to obtain annual responses. Panel (a) above corresponds to Panel (a) of Figure 3, while Panel (b) above corresponds to Panel (b) of Figure 3, and Panels (c) and (d) correspond to Panel (d) of Figure 3. Shaded bands indicate the range of outcomes from the lower bound estimate of $\sigma_{\omega,L} = 0.810$, to the upper bound estimate of $\sigma_{\omega,L} = \infty$, equivalent to the Full Segmentation case.

only vary the value of $\sigma_{\omega,L}$, and use the same estimates for the persistence and size of the shock across all three economies.

Panels (a) and (b) display results for the price-rent ratio and homeownership rate. Our estimation is successful, as the Benchmark model (red line) delivers a close fit of the empirical point estimates. The fit on the price-rent ratio is extremely close, while errors are slightly larger for the homeownership rate, reflecting the lower statistical precision around these estimates. In terms of slope, the Benchmark model delivers a response of the price-rent ratio between 6.98 and 9.31 times that of the homeownership rate depending on the horizon. Turning to polar models for comparison, the Full Segmentation model (blue line) delivers a nearly identical path of the price-rent ratio, showing that our Bench-

mark estimates imply a high degree of segmentation. However, the Full Segmentation model fails to deliver any increase in the homeownership rate. Last, the No Segmentation model delivers a much smaller rise in the price-rent ratio and a much larger rise in the homeownership rate, in line with the intuition from Section 2.

We next compute a “credible set” for $\sigma_{\omega,L}$ that reflects our 95% confidence interval for the inverse tenure supply slope estimates in Figure 3 Panel (d).²⁸ We choose values of $\sigma_{\omega,L}$ to minimize the distance to the upper and lower ends of the 95% confidence interval:

$$Q_{UB} = \sum_{k=2}^5 \left(\frac{IRF_{IR,k}^{Model} - (\beta_{IR,k} + 1.96 \times SE_{IR,k})}{SE_{v,k}} \right)^2$$

$$Q_{LB} = \sum_{k=2}^5 \left(\frac{IRF_{IR,k}^{Model} - (\beta_{IR,k} - 1.96 \times SE_{IR,k})}{SE_{v,k}} \right)^2,$$

where “IR” denotes the inverse slope ratio. For both re-estimations we hold the persistence of the shock constant at our prior estimate, while the shock size is irrelevant for the computation of this inverse ratio.

The resulting estimates are displayed as the shaded area in Figure 6 Panel (c). The upper bound, targeting the tops of the confidence intervals, yields an estimate of $\sigma_{\omega,L} = 0.810$, while the lower bound, targeting the bottoms of the confidence intervals, is most closely matched using the Full Segmentation case $\sigma_{\omega,L} \rightarrow \infty$. The figure shows that our Benchmark calibration provides a close fit of the inverse ratio as well.²⁹

To interpret the magnitudes implied by our credible set, Panel (d) displays the credible set and Benchmark responses for the inverse ratio alongside the responses from our No Segmentation and Full Segmentation economies. While the Full Segmentation economy forms the lower bound of our credible set by construction, the No Segmentation economy falls far outside of the credible set, with inverse ratios between 6 and 13 times those of our credible set upper bound, and between 4 and 32 times the upper bounds of our empirical confidence intervals. Our estimates can thus soundly reject the No Segmentation model.

²⁸We compute our credible set based on these ratio estimates, rather than the individual IRFs for two reasons. First, these estimates jointly summarize uncertainty about both the price-rent ratio response and the homeownership rate response into a single statistic at each horizon, which would otherwise be nontrivial to combine. Second, part of the standard errors for our individual IRFs represents uncertainty about the absolute scale of the shock rather than the *relative* size of the responses. Removing the size of the shock as a nuisance parameter allows for more statistical precision.

²⁹In principle, matching the separate price-rent ratio and homeownership rate IRFs in Figure 3 Panels (a) and (b) and matching the ratio responses in Panel (d) are both theoretically valid approaches. In practice, directly fitting the ratio responses delivers a slope so steep it is effectively identical to full segmentation. By matching the IRFs separately we are therefore opting for the more conservative set of estimates in terms of the influence of credit on house prices.

7 Model Results

Now that we have calibrated the model to match our empirical results, we run a series of experiments to quantitatively assess the role that credit played in the 2000s housing boom. We describe our experiments and results in detail below, and provide a summary table across all of our boom bust experiments in Table 2, as well as robustness of our main experiments to alternative values of $\sigma_{\omega,L}$ in Appendix Table A.1.

Table 2: Results, Boom Experiments

Experiment	Price-Rent	Homeown.	Loan-Inc.
Peak Data Increase	48.3%	3.3pp	72.2%
<i>Credit Relaxation (Share of Peak Data Increase)</i>			
Full Segmentation	38%	0%	53%
Benchmark	34%	27%	51%
Est. Lower Bound	26%	71%	46%
No Segmentation	-1%	201%	31%
<i>Credit Relaxation + Decline in Rates (Share of Peak Data Increase)</i>			
Full Segmentation	82%	0%	86%
Benchmark	72%	53%	80%
Est. Lower Bound	56%	135%	70%
No Segmentation	4%	353%	38%
<i>Removing Credit Relaxation from Full Boom (Share of Peak Data Increase)</i>			
Full Boom (Benchmark)	100%	100%	98%
No Credit Relaxation (Benchmark)	45%	59%	26%
Full Boom (No Segmentation)	100%	100%	103%
No Credit Relaxation (No Segmentation)	95%	-230%	50%
<i>Credit Relaxation + Decline in Rates: Extensions (Share of Peak Data Increase)</i>			
Landlord Credit, Not Recalibrated	81%	8%	80%
Landlord Credit, Recalibrated	80%	21%	80%
Saver Demand, Not Recalibrated	52%	63%	98%
Saver Demand, Recalibrated	54%	30%	102%

Notes: This table summarizes the results from the various nonlinear transition experiments in Sections 7 and 8. “Price-Rent” is the price-rent ratio, “Homeown.” is the homeownership rate, and “Loan-Inc.” is the aggregate loan to income ratio. The top row displays the actual changes in these variables, in levels from 1998:Q1 to the peak of each series during the boom period (2006 - 2008). The remaining numbers below display the shares of these peak increases explained by each model-experiment combination, calculated from 1998:Q1 to the peak of each model boom in 2007:Q1. The loan-to-income ratio is the ratio of household debt (FRED code: HMLBSHNO) to household gross income (FRED code: PI) in the Flow of Funds. For other data definitions see the notes for Figure 1.

To begin, we simulate a realistic relaxation of credit standards and evaluate the model's implications for the evolution of debt and house prices. Our baseline experiment follows KMV in relaxing LTV limits from 85% to 99% and PTI limits from 36% to 65% unexpectedly and permanently 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 computed as nonlinear perfect foresight paths.

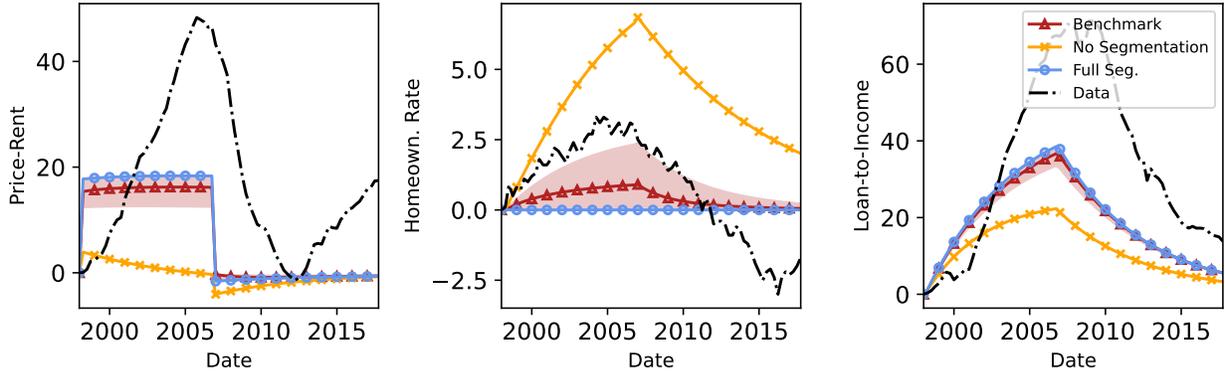
The results of this experiment are shown visually in Figure 7. To highlight the role of landlord heterogeneity, we again plot the responses in our Benchmark model against the polar No Segmentation (perfectly elastic supply) and Full Segmentation (perfectly inelastic supply) alternatives. As in Figure 6, the shaded bands account for the credible set for $\sigma_{\omega,L}$. Our Benchmark model displays a large price response to the credit standard shifts, accounting for 34% of the peak rise in price-rent ratios observed in the boom, while a model setting $\sigma_{\omega,L}$ to the lower bound of the credible set would explain 26%. This stands in sharp contrast to the No Segmentation model, where the same credit relaxation explains -1% of the peak growth in price-rent ratios, as landlords completely satisfy the increase in demand, preventing a rise in prices. Instead, 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-rent ratios.

This finding for house prices also has important implications for credit growth. While credit standards are loosened equally for all three cases depicted in Figure 7, credit growth over the boom is much larger in the Benchmark economy relative to the No Segmentation economy, explaining 51% and 31% of the observed rise, respectively. 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. Consequently, the same credit loosening leads to much more levered households in the Benchmark economy after the reversal.

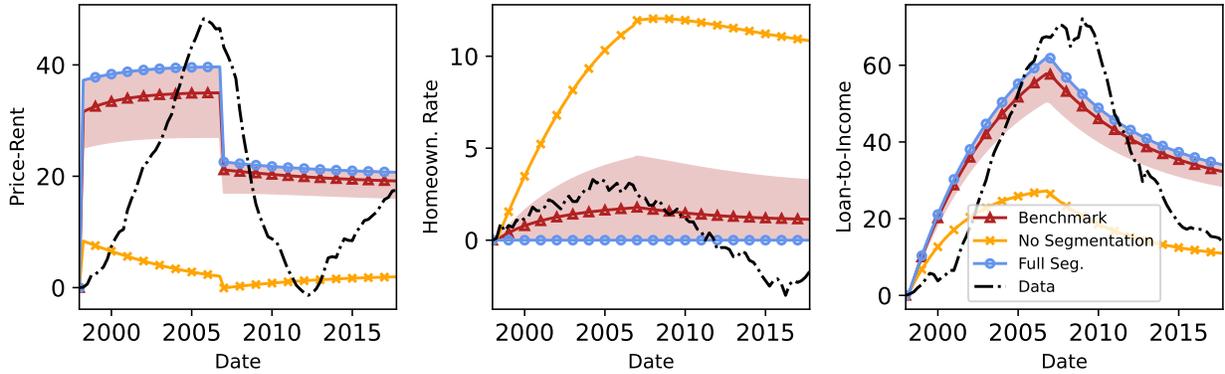
For a more comprehensive view of the role of credit, we next incorporate an additional 2ppt fall in mortgage spreads, assumed to be permanent, which reflects secular declines in interest rates over the boom period. This causes an outward shift of housing demand, which, given our estimated rental frictions generates a large additional increase in house prices. Combining the relaxation in LTV and PTI limits and the fall in rates can explain 72% of the observed rise in price-rent ratios and 80% of the rise in loan-to-income ratios, shown in Figure 7 Panel (b). These results again stand in contrast to the 4% and 38% shares explained in the No Segmentation model, in which neither the price nor quantity of credit is an important determinant of the price-rent ratio.

Beyond changes in credit prices and standards, non-credit factors such as overopti-

Figure 7: Credit Relaxation Experiment



(a) Credit Standards Only



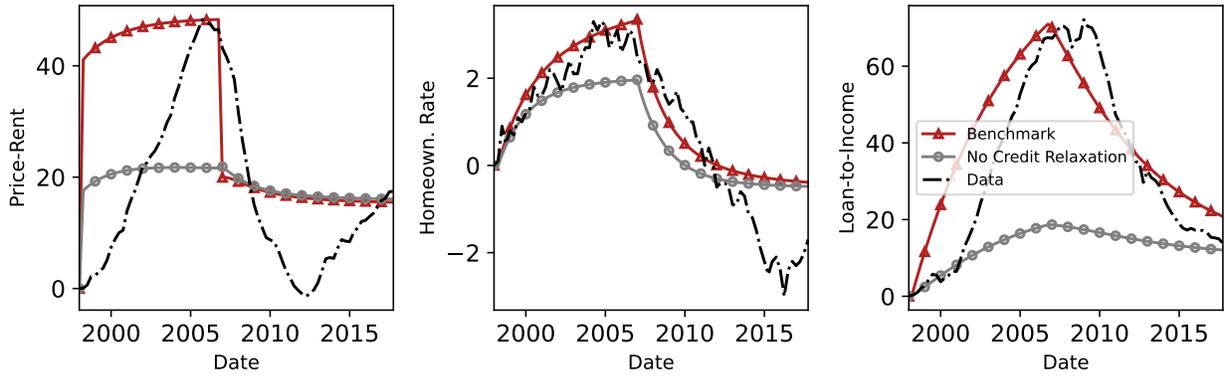
(b) Credit Standards + Falling Rates

Notes: Each panel displays perfect foresight paths following a relaxation LTV and PTI constraints. The “Benchmark” model sets a value of $\sigma_{\omega,L}$ calibrated to match our empirical IRFs as in Section 6, while the “No Segmentation” model sets $\sigma_{\omega,L} = 0$ and the “Full Segmentation” model sets $\sigma_{\omega,L} \rightarrow \infty$. Shaded bands indicate the range of outcomes from the lower bound estimate of $\sigma_{\omega,L} = 0.810$, to the upper bound estimate of $\sigma_{\omega,L} = \infty$, equivalent to the Full Segmentation case. Results are summarized numerically in Table 2. For data definitions see notes for Figure 1 and Table 2.

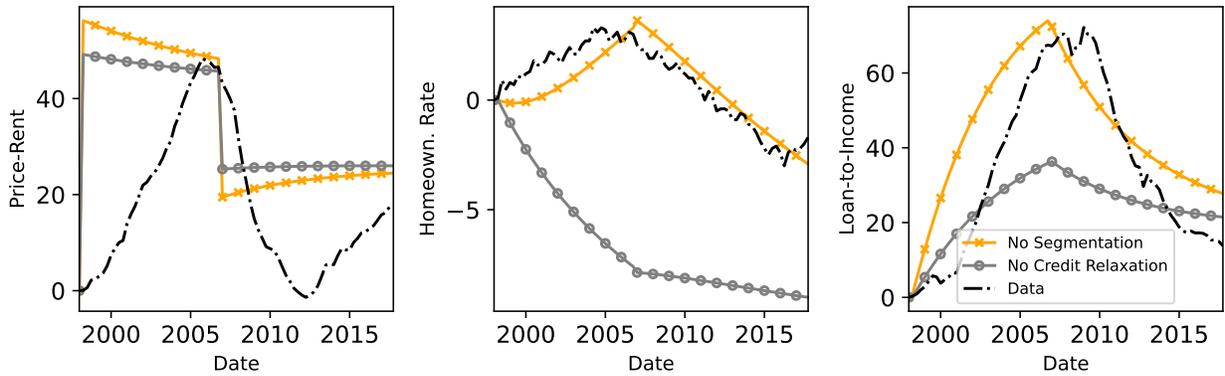
mistic house price expectations are also widely believed to have played a major role in driving the boom (see e.g., Kaplan et al. (2020)). Since strong interactions can be present between credit conditions and these non-credit factors (Greenwald (2018)), we now incorporate these residual non-credit factors on top of the relaxation in LTV and PTI constraints and 2% decline in mortgage rates to create a “Full Boom” experiment.

We follow the intuition in Section 2 that for any credit shock and supply curve slope we can use additional shifts to the demand and tenure supply curves to exactly match the total increase in both the price-rent ratio and in the homeownership rate over the complete boom period. We implement these shocks as level shifts in the ownership utility

Figure 8: Full Boom Experiment



(a) Benchmark Economy



(b) No Segmentation Economy

Notes: Plots display perfect foresight paths. The “Benchmark” calibrates $\sigma_{\omega,L}$ to match our empirical IRFs as in Section 6, while the “No Segmentation” model sets $\sigma_{\omega,L} = 0$. For each set of plots, the colored plots display an experiment imposing a credit relaxation, decline in the interest rate, and level shifts to the demand and supply curves ($\mu_{\omega,B}$ and $\mu_{\omega,L}$), with these level shifts chosen to exactly match the peak growth of price-rent ratios and the homeownership rate over the housing boom. In each panel, the “No Credit Relaxation” responses display an alternative experiment showing the same decline in the mortgage rate, and level shifts to our demand and supply curves while removing the relaxation of credit standards. Results are summarized numerically in Table 2. For data definitions see notes for Figure 1 and Table 2.

distributions through changes in $\mu_{\omega,B}$ and $\mu_{\omega,L}$, respectively, that are assumed permanent during the boom, then revert to their original values in the bust.³⁰ To complete this “Full Boom” experiment, we incorporate additional features relevant to the bust: a further 3ppt fall in both mortgage rates and the landlord discount rate, consistent with a broad decline in long-term interest rates, and a 10% tightening of both LTV and PTI limits, consistent with a further tightening of credit standards.³¹

³⁰This parsimonious shock specification avoids taking a stand on whether borrowers and landlords share symmetric beliefs, and avoids the need for additional shocks to explain the aggregate data.

³¹This is best interpreted as tight credit score standards preventing a fraction of the population from

The resulting transition paths are plotted in Figure 8, Panel (a). Overall, these assumptions generate a reasonably good fit of the dynamics of the boom and bust, with two main exceptions: (i) house prices jump in the model rather than adjust sluggishly in the data, which is typical for models lacking search frictions; and (ii) our model “bust” features a softer landing relative to the data, as we lack the foreclosures and financial market features that transformed the housing crash into a global financial crisis.

To measure the total contribution of credit conditions in this simulated boom-bust, we then remove the simulated credit expansion, while leaving all the other factors in place, to generate the series labeled “No Credit Relaxation.” We find that removing the credit expansion from our Full Boom experiment would have reduced the overall rise in price-rent ratios by 55% and in loan-to-income ratios by 74% in our Benchmark economy. These shares, 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 (34% and 51%, respectively), because loose credit amplifies the role of non-credit demand factors. The simple intuition, discussed at length in Greenwald (2018), is that even if borrower households perceive large gains to ownership, they lack the financial resources to pay for large fractions of their housing purchases in cash. When expectations rise without a concurrent relaxation in credit standards, binding PTI limits constrain households’ ability to finance these properties, dampening the growth of house price and credit volumes.

This final set of results indicates that macroprudential policy that restricts credit through LTV and PTI limits is effective at restraining a housing boom.³² This finding depends heavily on our key parameter $\sigma_{\omega,L}$ and the slope of the supply curve. To show this, Figure 8 Panel (b) replicates the exact same experiment using the frictionless No Segmentation model. As before, we add a set of demand and supply shocks to our relaxation of credit and fall in rates to exactly recreate the entire boom in price-rent ratios and homeownership, and then solve a second transition that removes the relaxation of credit standards. In the absence of rental market frictions, removing credit standard relaxation reduces the increase in price-rent ratios by only 5%. Because it fails to stem the rise in collateral values, this tight credit counterfactual is also much less effective at reducing credit growth, with the rise in loan-to-income ratios reduced by only 50%. The fact that alternate calibrations can fully explain the rise in price-rent ratios and homeownership rates while having strikingly different implications for the effectiveness of macroprudential policy illustrates the importance of using the tenure supply curve slope to discipline models.

obtaining credit at all, rather than a decline in maximum LTV and PTI ratios at the intensive margin.

³²Greenwald (2018) shows that among these, PTI limits are much more effective than LTV limits at dampening house price growth in booms.

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. A relaxation of credit standards can explain roughly between one third and one half of the rise in price-rent ratios, depending on the order it is added relative to other shocks — results that are much closer to the extreme of full segmentation than to a frictionless model with no landlord heterogeneity.

8 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 heterogeneous costs in converting properties between owner-occupied and renter-occupied. While this assumption makes the economics of the model transparent, it is clearly an abstraction, as many landlords also face financial constraints in reality. As discussed earlier, we have also abstracted from trade in housing between borrowers and savers. In this section, we extend the model to relax each of these assumptions in turn.

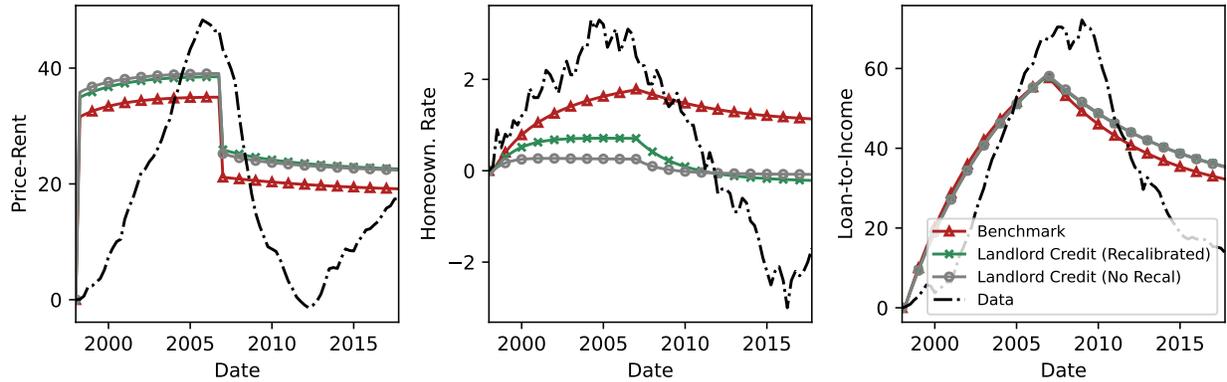
8.1 Landlord Credit

In practice, landlords are not deep-pocketed, and the vast majority of investor-owned properties are purchased with mortgages. To capture this, we now allow landlords to purchase properties with mortgage credit that is affected by changes in credit supply.

We begin by reassessing the intuition developed in the supply and demand framework of Section 2. Recall that a credit relaxation shifts the demand curve but not the tenure supply curve, causing movement in the price-rent ratio, the homeownership rate, or both, as the equilibrium travels up the supply curve in Figure 2 Panel (c). Introducing credit for landlords implies that a relaxation in credit not only shifts the demand curve upward but also shifts the tenure supply curve upward, as in Figure 2 Panel (d). Adding landlord credit to the baseline model while holding the parameters fixed will lead to a smaller (or potentially even negative) change in the homeownership rate and a larger change in the price-rent ratio. This is represented by a shift from the solid supply curve to the dashed supply curve in Figure 2 Panel (d).

To illustrate this intuition quantitatively, we implement a version of the model with landlord credit, with a full description in Appendix A.2. We assume that landlords use a parallel borrowing technology to borrowers, with an LTV limit of 65% (a standard constraint for multi-family construction loans) and no PTI limit.

Figure 9: Credit Standards + Falling Rates Experiment, Landlord Credit Extension



Notes: Plots display perfect foresight paths following a relaxation of credit standards and a decline in interest rates. The “Benchmark” model sets a value of $\sigma_{\omega,L}$ calibrated to match our empirical IRFs as in Section 6. The “Landlord Credit (No Recal)” model applies the landlord credit extension holding $\sigma_{\omega,L}$ fixed as in our Benchmark calibration, while the “Landlord Credit (Recalibrated)” model applies the same extension while recalibrating $\sigma_{\omega,L}$ under the new model. Results are summarized numerically in Table 2. For data definitions see notes for Figure 1 and Table 2.

We first solve this extension using our Benchmark value of $\sigma_{\omega,L}$. Figure 9 displays the results from an experiment analogous to that of Figure 7 Panel (b), which both relaxes credit conditions and allows interest rates to fall, with summary statistics again displayed in Table 2. To provide a quantitative example of a loosening of landlord credit, we assume that landlord mortgages face an equal decline in rates and that landlord credit also expands to a new LTV limit of 85% during the boom.

The resulting responses show that, holding parameters fixed (the path denoted “Landlord Credit (No Recal)”), adding landlord credit increases the response of the price-rent ratio, explaining 81% of the rise observed in the data, compared to 72% for the Benchmark model. At the same time, the landlord credit model features a smaller rise in the homeownership rate, explaining only 8% of the rise in the data, compared to 53% for the Benchmark model. These results are consistent with the intuition in Figure 2 Panel (d).

The results holding parameters fixed would, however, make the model inconsistent with our empirical regressions.³³ To address this, we repeat our exercise in Section 6 to recalibrate $\sigma_{\omega,L}$ for the landlord credit model.³⁴ The resulting responses, denoted “Landlord Credit (Recalibrated),” follow a very similar pattern, explaining 80% of the observed

³³We note that under this extension the ratio estimated by our regressions now reflects a locus of equilibria as both demand and supply shift, rather than the slope of the supply curve alone.

³⁴This recalibration requires an adjustment as roughly half of rental units are located in multifamily buildings too large to be affected by the changes in the CLL on which the LS instrument is based (see Appendix A.2 for details).

rise in the price-rent ratio, and 21% of the rise in the homeownership rate.

Overall, these results indicate that incorporating landlord credit and its relaxation during the housing boom period would strengthen the role of credit in driving house prices. As a result, we believe that our Benchmark calibration is, if anything, conservative, and should provide a lower bound on the true contribution of credit over this period.

8.2 Saver Housing Demand

Our baseline model also assumes that housing demand by unconstrained households (“savers”) is fixed. As shown in Justiniano et al. (2015), Kaplan et al. (2020), and Kiyotaki et al. (2011), these savers have a relatively constant marginal utility and are not credit constrained at the margin so their trade in housing absorbs demand by constrained borrower households and dampens the impact of credit on prices.

To embed this mechanism in our model, we relax our assumption of fixed saver demand $H_{s,t} = \bar{H}_s$ and allow savers to freely trade housing. This implies the additional optimality condition from the saver first order condition:

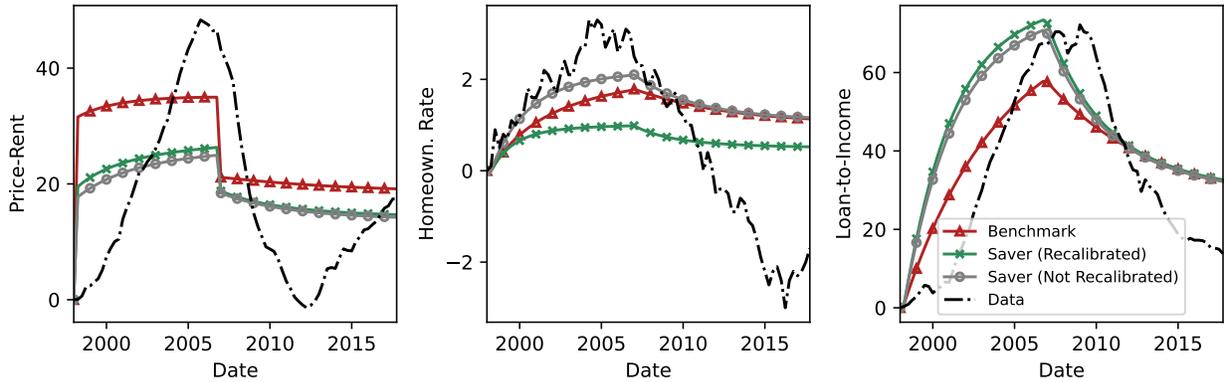
$$p_t^{\text{Saver}} = E_t \left\{ \Lambda_{t+1}^S \left[\underbrace{u_{h,t}^S / u_{c,t}^S}_{\text{housing services}} + \underbrace{(1 - \delta) p_{t+1}}_{\text{continuation value}} \right] \right\}. \quad (9)$$

This expression is nearly identical to the borrower’s condition (6) with two exceptions. First, the collateral value term \mathcal{C} is equal to zero, as the saver does not use credit. Second, we assume no saver heterogeneity ($\omega_{i,t}^S = 0$). Instead, reaching an equilibrium where $p^{\text{Saver}} = p^{\text{Demand}} = p^{\text{Supply}}$ occurs entirely through changes in saver housing H_S , which adjusts the marginal utility term $u_{h,t}^S / u_{c,t}^S$. Since heterogeneity would steepen the slope of the saver demand curve, diminishing their ability to absorb changes in borrower demand, these results represent an upper bound on the role of savers.

Figure 10 compares the response to our experiment in which credit standards are loosened and interest rates fall between our Benchmark calibration and this saver demand extension. As before, we plot one version holding $\sigma_{\omega,L}$ fixed (“Not Recalibrated”) and a second version (“Recalibrated”) after repeating the $\sigma_{\omega,L}$ calibration procedure in Section 6. Beginning with the non-recalibrated response, we observe that the rise in price-rent ratios is diminished as savers react to the rise in prices by selling portions of their housing stock to borrowers, absorbing demand. Since these savers are still homeowners, there is no major change in the response of the homeownership rate.

However, introducing savers while holding $\sigma_{\omega,L}$ fixed worsens the model’s fit of our

Figure 10: Credit Standards + Falling Rates Experiment, Saver Demand Extension



Notes: Plots display perfect foresight paths following a relaxation of credit standards and a decline in interest rates. The “Benchmark” model sets a value of $\sigma_{\omega,L}$ calibrated to match our empirical IRFs as in Section 6. The “Saver (Not Recalibrated)” model applies the flexible saver demand extension holding $\sigma_{\omega,L}$ fixed as in our Benchmark calibration, while the “Saver (Recalibrated)” model applies the same extension while recalibrating $\sigma_{\omega,L}$ under the new model. Results are summarized numerically in Table 2. For data definitions see notes for Figure 1 and Table 2.

empirical IRFs in Section 3.2, as the price-rent ratio increases by too little relative to the homeownership rate. Recalibrating $\sigma_{\omega,L}$ to restore this fit yields the “Recalibrated” response, which yields a slightly larger rise in price-rent ratios and a much smaller change in the homeownership rate, effectively restoring the correct ratio.³⁵ Even with a perfectly frictionless saver margin, the recalibrated saver model still explains 54% of the observed rise in the price-rent ratio from changes in the price and quantity of credit alone. While this response is about 25% smaller than the 72% observed in the Benchmark model, it does not overturn our core results.

We consider this saver extension to be an extreme lower bound on the strength of credit on house prices. While savers in our model are able to frictionlessly adjust the size of their home at the intensive margin in response to the housing cycle, housing is in reality both indivisible and highly heterogeneous in both location and quality. In practice, it is not a viable option for saver households to sell portions of their homes to borrowers when credit relaxes and rebuy these portions when credit tightens. Instead, Landvoigt et al. (2015) show that while changes in demand can ripple up or down the housing quality ladder, this effect is still significantly muted relative to a frictionless benchmark, implying that the real world likely falls closer to our benchmark model than our saver extension.

³⁵The reason the recalibration ends up mostly adjusting along the homeownership margin rather than the price-rent margin is that the price-rent ratio response in the Benchmark model is already very close to the Full Segmentation model, leaving little room for further increases as $\sigma_{\omega,L}$ rises.

9 Conclusion

More than a decade after the Great Recession there is still a lack of consensus about the role of credit supply in explaining house prices dynamics over the boom and bust. We argue this is because most of the literature has focused on two polar assumptions for the degree of segmentation in housing markets between credit-sensitive borrowers and credit-insensitive agents such as landlords or unconstrained savers.

In this paper, we generalize these polar cases to allow for arbitrary intermediate levels of rental frictions. Building on supply-demand intuition, we show that the causal effect of credit on the price-rent ratio relative to the effect of the same shock on the homeownership rate is a sufficient statistic for determining the degree of frictions. Using three sets of instrumental variables, we show that credit supply shocks cause a significant increase in price-rent ratios and a more muted and statistically insignificant homeownership response. Calibrating a model to match these estimates, we find that credit supply can explain between 35% and 54% of the rise in price-rent ratios over the 2000s housing boom. 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, which are often overlooked, but are central to many core results. We hope that our findings motivate future work that both uses and develops intermediate models in place of either polar assumption and that refines the tenure supply slope estimates we propose researchers use to calibrate such models.

References

- ADELINO, M., A. SCHOAR, AND F. SEVERINO (2020): "Credit Supply and House Prices: Evidence From Mortgage Market Segmentation," Tech. rep., National Bureau of Economic Research.
- BERGER, D., N. TURNER, AND E. ZWICK (2020): "Stimulating Housing Markets," *The Journal of Finance*, 75, 277–321.
- CAMPBELL, J. R. AND Z. HERCOWITZ (2005): "The Role of Collateralized Household Debt in Macroeconomic Stabilization," Tech. rep., National Bureau of Economic Research.
- CHEN, V. W. (2019): "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.
- EGGERTSSON, G. B. AND P. KRUGMAN (2012): "Debt, Deleveraging, and the Liquidity Trap: A Fisher-Minsky-Koo Approach," *The Quarterly Journal of Economics*, 127, 1469–1513.
- 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.
- FOOTE, C. L., K. S. GERARDI, AND P. S. WILLEN (2012): "Why Did So Many People Make So Many Ex Post Bad Decisions? The Causes of the Foreclosure Crisis," Tech. rep., National Bureau of Economic Research.
- GARRIGA, C. AND A. HEDLUND (2018): "Housing Finance, Boom-Bust Episodes, and Macroeconomic Fragility," Working Paper, Federal Reserve Bank of St. Louis and University of Missouri.
- (2020): "Mortgage Debt, Consumption, and Illiquid Housing Markets in the Great Recession," *American Economic Review*, 110, 1603–1634.
- GARRIGA, C., F. E. KYDLAND, AND R. ŠUSTEK (2017): "Mortgages and monetary policy," *The Review of Financial Studies*, 30, 3337–3375.
- GARRIGA, C., R. MANUELLI, AND A. PERALTA-ALVA (2019): "A Macroeconomic Model of Price Swings in the Housing Market," *American Economic Review*, 109, 2036–2072.
- GETE, P. AND M. REHER (2018): "Mortgage Supply and Housing Rents," *Review of Financial Studies*, 31, 4884–4911.
- GHENT, A. (2012): "Infrequent Housing Adjustment, Limited Participation, and Monetary Policy," *Journal of Money, Credit and Banking*, 44, 931–955.
- 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.
- 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 (2021): "Mortgage design in an equilibrium model of the housing market," *The Journal of Finance*, 76, 113–168.
- HALKET, J., L. NESHEIM, AND F. OSWALD (2020): "The housing stock, housing prices, and user costs: The roles of location, structure, and unobserved quality," *International Economic Review*, 61, 1777–1814.

- HUO, Z. AND J.-V. RIOS-RULL (2016): "Financial Frictions, Asset Prices, and the Great Recession," Working Paper, Federal Reserve Bank of Minneapolis.
- IACOVIELLO, M. (2005): "House Prices, Borrowing Constraints, and Monetary Policy in the Business Cycle," *American Economic Review*, 739–764.
- IACOVIELLO, M. AND S. NERI (2010): "Housing Market Spillovers: Evidence From an Estimated DSGE Model," *American Economic Journal: Macroeconomics*, 2, 125–164.
- JOINT CENTER FOR HOUSING STUDIES OF HARVARD UNIVERSITY (2020): "The State of the Nation's Housing 2020," Tech. rep.
- 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 (2015): "Household Leveraging and Deleveraging," *Review of Economic Dynamics*, 18, 3–20.
- (2017): "The Mortgage Rate Conundrum," *Journal of Political Economy*, forthcoming.
- (2019): "Credit Supply and the Housing Boom," *Journal of Political Economy*, 127, 1317–1350.
- KAPLAN, G., K. MITMAN, AND G. L. VIOLANTE (2020): "The housing boom and bust: Model meets evidence," *Journal of Political Economy*, 128, 3285–3345.
- 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 (2013): "Land-Price Dynamics and Macroeconomic Fluctuations," *Econometrica*, 81, 1147–1184.
- (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 (2021): "Credit Supply and Housing Speculation," Forthcoming.
- MONACELLI, T. (2008): "Optimal Monetary Policy With Collateralized Household Debt and Borrowing Constraints," in *Asset Prices and Monetary Policy*, University of Chicago Press, 103–146.
- NADAULD, T. D. AND S. M. SHERLUND (2013): "The Impact of Securitization on the Expansion of Subprime Credit," *Journal of Financial Economics*, 107, 454–476.
- 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.

ROGNLIE, M., A. SHLEIFER, AND A. SIMSEK (2018): "Investment hangover and the great recession," *American Economic Journal: Macroeconomics*, 10, 113–53.

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 Model Appendix

A.1 Equilibrium Conditions

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

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

$$(h_{B,t}) : \quad q_t = (u_{B,t}^h / u_{B,t}^c) \quad (10)$$

$$(H_{B,t}^*) : \quad p_t = \frac{E_t \left\{ \Lambda_{B,t+1} \left[\bar{\omega}_{B,t} + q_{t+1} + \left(1 - \delta - (1 - \rho_B) \mathcal{C}_{B,t+1} \right) p_{t+1} \right] \right\}}{1 - \mathcal{C}_{B,t}} \quad (11)$$

$$(M_{B,t}^*) : \quad 1 = \Omega_{M,t}^B + r_{j,t}^* \Omega_{X,t}^B + \mu_{B,t}, \quad (12)$$

where:

$$\mathcal{C}_{B,t} = \mu_{B,t} F_{B,t}^{LTV} \theta_{B,t}^{LTV}.$$

$\mu_{B,t}$ is the multiplier on the borrowing constraint, $F_{B,t}^{LTV}$ 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 $\Omega_{M,t}^B$ and of interest payments $\Omega_{X,t}^B$ satisfy:

$$\begin{aligned} \Omega_{M,t}^B &= E_t \left\{ \Lambda_{B,t+1} \tau^{-1} \left[\nu_B + (1 - \nu_B) \left(\rho_B + (1 - \rho_B) \Omega_{M,t+1}^B \right) \right] \right\} \\ \Omega_{X,t}^B &= E_t \left\{ \Lambda_{B,t+1} \tau^{-1} \left[(1 - \tau) + (1 - \nu_B) (1 - \rho_B) \Omega_{X,t+1}^B \right] \right\}. \end{aligned}$$

Equation (10) sets the rent equal to the marginal rate of substitution between housing services and consumption. Equation (11) 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 $(1 - \rho_B)$. Equation (12) 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).

Landlord's Problem. The landlord's optimality conditions is:

$$(H_{B,t}^*) : \quad p_t = \frac{E_t \left\{ \Lambda_{L,t+1} \left[\bar{\omega}_{L,t} + q_{t+1} + \left(1 - \delta - (1 - \rho_{L,t+1}) \mathcal{C}_{L,t+1} \right) p_{t+1} \right] \right\}}{1 - \mathcal{C}_{L,t}}.$$

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

$$(B_t) : \quad 1 = R_t E_t \left[\pi^{-1} \Lambda_{S,t+1} \right]$$

$$(M_{B,t}^*) : \quad 1 = Q_{M,t}^S + r_{B,t}^* Q_{X,t}^S,$$

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

$$Q_{M,t}^S = E_t \left\{ \Lambda_{S,t+1} \pi^{-1} \left[v_B + (1 - v_B) \left(\rho_B + (1 - \rho_B) Q_{M,t+1}^S \right) \right] \right\}$$

$$Q_{X,t}^S = E_t \left\{ \Lambda_{S,t+1} \pi^{-1} \left[(1 - \tau) + (1 - v_B) (1 - \rho_B) Q_{X,t+1}^S \right] \right\}.$$

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

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

$$1 = p_t (1 - \varphi) L_t^\varphi Z_t^{-\varphi}.$$

A.2 Extension: Landlord Credit

When landlords use credit, the landlord's budget constraint becomes:

$$c_{L,t} \leq \underbrace{(1 - \tau) y_{L,t}}_{\text{after-tax income}} + \underbrace{\rho_L \left(M_{L,t}^* - \pi^{-1} (1 - v_L) M_{L,t-1} \right)}_{\text{net mortgage iss.}} - \underbrace{\pi^{-1} (1 - \tau) X_{L,t-1}}_{\text{interest payment}} - \underbrace{v_L \pi^{-1} M_{L,t-1}}_{\text{principal payment}}$$

$$- \underbrace{\rho_L p_t (H_{L,t}^* - H_{L,t-1})}_{\text{net housing purchases}} - \underbrace{\delta p_t H_{L,t-1}}_{\text{maintenance}} + \underbrace{q_t H_{L,t-1}}_{\text{rent}} + \underbrace{\left(\int_{\bar{\omega}_{L,t-1}} \omega d\Gamma_{\omega,L} \right) \hat{H}_{t-1}}_{\text{owner surplus}} + \underbrace{T_{L,t}}_{\text{rebates}},$$

while the landlord's laws of motion are:

$$M_{L,t} = \rho_L M_{L,t}^* + (1 - \rho_L) (1 - v_L) \pi^{-1} M_{L,t-1}$$

$$X_{L,t} = \rho_L r_{L,t}^* M_{L,t}^* + (1 - \rho_L) (1 - v_L) \pi^{-1} X_{L,t-1}$$

$$H_{L,t} = \rho_L H_{L,t}^* + (1 - \rho_L) H_{L,t-1}.$$

We assume that the landlord also faces the LTV limit:

$$M_{L,t}^* \leq \theta_L^{LTV} p_t H_{L,t}^*$$

The landlord's optimality conditions are:

$$(H_{B,t}^*) : \quad p_t = \frac{E_t \left\{ \Lambda_{L,t+1} \left[\bar{\omega}_{L,t} + q_{t+1} + \left(1 - \delta - (1 - \rho_{L,t+1}) \mathcal{C}_{L,t+1} \right) p_{t+1} \right] \right\}}{1 - \mathcal{C}_{L,t}}$$

$$(M_{L,t}^*) : \quad 1 = \Omega_{M,t}^L + r_{j,t}^* \Omega_{X,t}^L + \mu_{L,t}$$

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

$$\Omega_{M,t}^L = E_t \left\{ \Lambda_{L,t+1} \pi^{-1} \left[v_L + (1 - v_L) \left(\rho_{L,t+1} + (1 - \rho_{L,t+1}) \Omega_{M,t+1}^L \right) \right] \right\}$$

$$\Omega_{X,t}^L = E_t \left\{ \Lambda_{L,t+1} \pi^{-1} \left[(1 - \tau) + (1 - v_L) (1 - \rho_{L,t+1}) \Omega_{X,t+1}^L \right] \right\},$$

symmetric to the borrower case.

The saver's budget constraint becomes:

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

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

where the s_j s control the degree to which spreads react to the spread shock Δ_t , used in our recalibration exercise below.

Calibration. We assume that landlords face a 65% LTV limit, and no PTI limit, corresponding to the parameter values $\theta_L^{LTV} = 0.65$ and $\theta_L^{PTI} = \infty$, and implying $F_{L,t}^{LTV} = 1$. To calibrate the model to match our empirical IRFs as in Section 6, we need to map our identified LS shock into the model. While this mapping into a shock to mortgage spreads is straightforward for the model with borrower credit only, as the vast majority of conforming mortgages are securitized by the GSEs, this is not the case for landlords. While single family rental properties, and multifamily rental properties with 2-4 units, can be financed by GSE loans, and are thus influenced by changes in the conforming loan limit,

rental properties with more than four units are not eligible for this type of mortgage, and thus should not be directly affected by the LS shock. Since roughly 50% of rental units are in eligible 1-4 family buildings (Joint Center for Housing Studies of Harvard University (2020)), we choose the parsimonious recalibration $s_B = 1, s_L = 0.5$, so that half of landlord credit is eligible for the GSE subsidy.

A.3 Model: Additional Results

This section presents additional model results referenced in the main text. To begin, Table A.1 reproduces our main credit relaxation experiments for various values of $\sigma_{\omega,L}$. For easier interpretation, and to connect back to the implied ratios we computed in Section 4, we recalibrate our model for a set of ratios between 1 and 10 so that our model implied IRFs, computed as in Figure 6, so that the response of the price-rent ratio to the response of the homeownership rate at the two-year horizon is exactly equal to this ratio.

Figure A.1 shows results for our Benchmark model varying the borrower heterogeneity parameter $\sigma_{\omega,B}$. Panel (a) displays results for the “Credit Relaxation” experiment from Figure 7a, while Panel (b) displays results for the “Credit Relaxation + Rates” experiment from Figure 7b. Since our borrower heterogeneity parameter $\sigma_{\omega,B}$ is calibrated to match the number of rent to own switches in a hypothetical First Time Homebuyer Credit experiment, the “Higher Dispersion” series targets a number of switchers half as large as in our Benchmark calibration, while the “Lower Dispersion” series targets a number of switchers twice as large as in our Benchmark calibration. For intuition, higher dispersion means that borrowers differ more in their valuations of housing, meaning that fewer households need to switch to adjust the marginal buyer’s valuation and clear the market. For both alternative models, we do not recalibrate $\sigma_{\omega,L}$. Figure A.1 shows that the response series are virtually identical, reinforcing that borrower dispersion is not a particularly important parameter for our results for any value within the reasonable range.

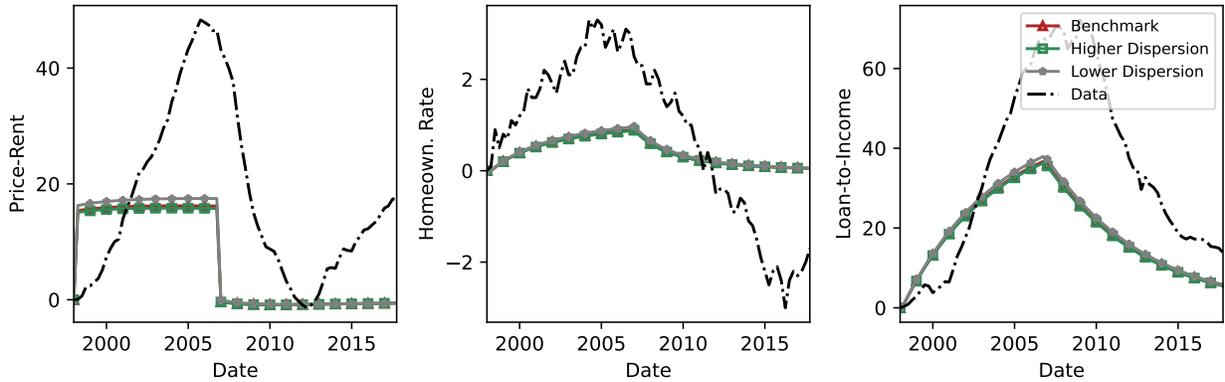
Last, Figure A.2 compares results under our Benchmark model to those from an alternative model with fixed housing supply ($H_t = \bar{H}$ for all t). The figure shows that the responses in the two economies are largely similar, with the fixed-supply model producing a slightly larger price-rent ratio response. The intuition behind this finding is that, while construction supply affects the degree to which housing demand influences house prices or rents, it has a much smaller impact on the ratio of prices to rents, which is the key object we study. This can be seen in the second row of Figure A.2, where we see a larger response of both prices and rents in the model with fixed construction supply.

Table A.1: Results, Boom Experiments, by Target Ratio

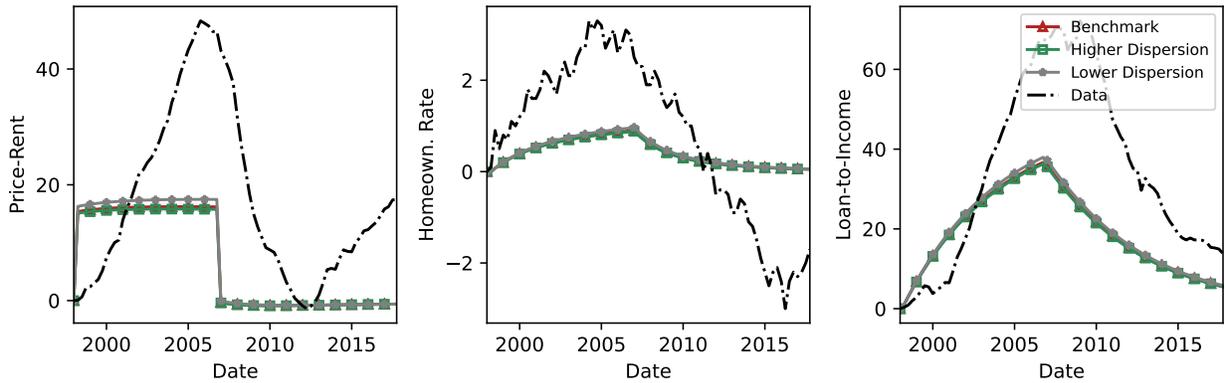
Experiment	Price-Rent	Homeown.	Loan-Inc.
Peak Data Increase	48.3%	3.3pp	72.2%
<i>Credit Relaxation (Share of Peak Data Increase)</i>			
Ratio = 1	5%	175%	34%
Ratio = 2	16%	122%	41%
Ratio = 3	22%	89%	44%
Ratio = 4	26%	68%	47%
Ratio = 5	29%	54%	48%
Ratio = 6	31%	44%	49%
Ratio = 7	32%	37%	50%
Ratio = 8	33%	32%	50%
Ratio = 9	33%	28%	51%
Ratio = 10	34%	25%	51%
<i>Credit Relaxation + Decline in Rates (Share of Peak Data Increase)</i>			
Ratio = 1	15%	313%	44%
Ratio = 2	36%	226%	58%
Ratio = 3	49%	168%	66%
Ratio = 4	57%	130%	71%
Ratio = 5	62%	105%	74%
Ratio = 6	66%	86%	76%
Ratio = 7	69%	73%	78%
Ratio = 8	71%	63%	79%
Ratio = 9	72%	55%	80%
Ratio = 10	73%	49%	81%

Notes: This table displays results varying $\sigma_{\omega,L}$ for the “Credit Relaxation” and “Credit Relaxation + Decline in Rates” experiments from Figure 7 in Section 7. Each row corresponds to a calibration of $\sigma_{\omega,L}$ chosen so that for e.g., Ratio = 5, our model-implied IRFs computed as in Figure 6 have a price-rent response that is 5 times larger than the homeownership rate response at the 2-year horizon. “Price-Rent” is the price-rent ratio, “Homeown.” is the homeownership rate, and “Loan-Inc.” is the aggregate loan to income ratio. The top row displays the actual changes in these variables, in levels from 1998:Q1 to the peak of each series during the boom period (2006 - 2008). The remaining numbers below display the shares of these peak increases explained by each model-experiment combination, calculated from 1998:Q1 to the peak of each model boom in 2007:Q1. For data definitions see notes for Figure 1 and Table 2.

Figure A.1: Credit Relaxation by Borrower Heterogeneity



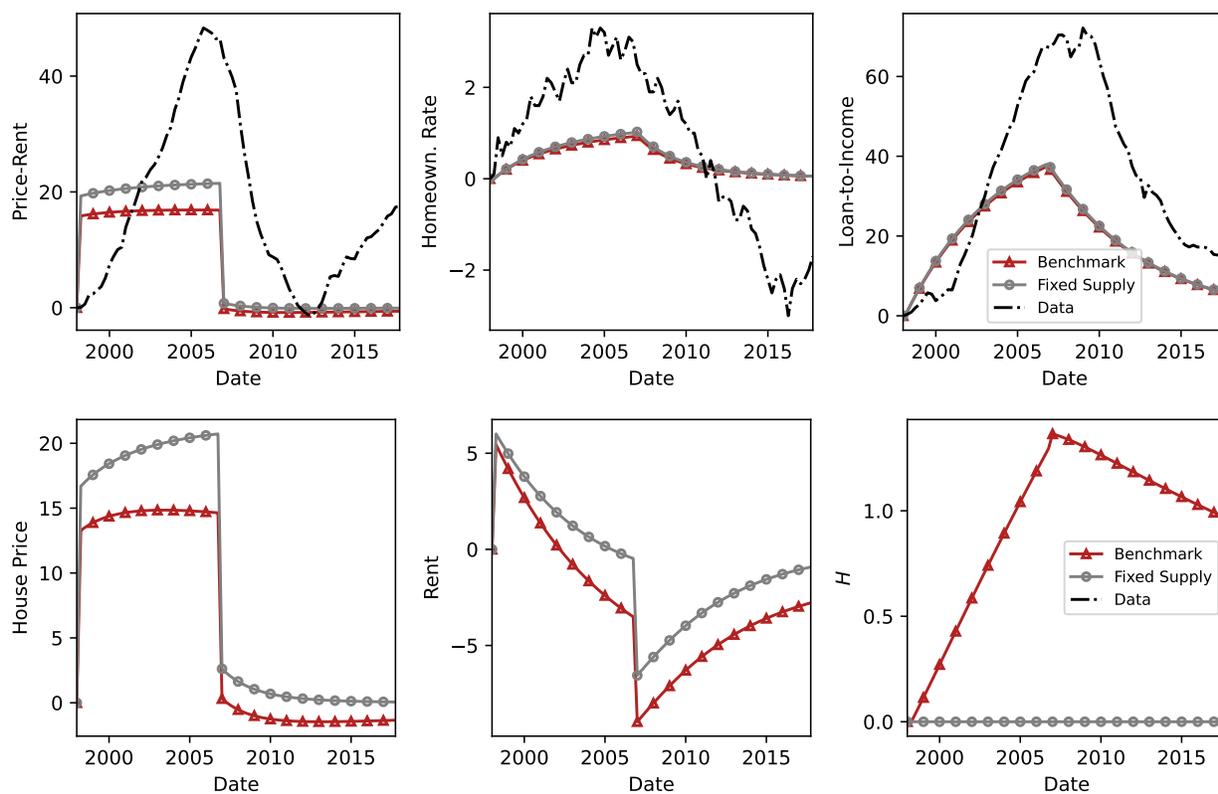
(a) Credit Relaxation by Borrower Heterogeneity



(b) Credit Relaxation by Borrower Heterogeneity

Notes: Panel (a) shows responses to our Credit Relaxation experiment from Figure 7a varying the level of borrower dispersion (σ_B). Panel (b) repeats this exercise for our Credit Relaxation + Rates experiment from Figure 7b. The “Higher Dispersion” series sets $\sigma_{\omega,B}$ so that half as many renters ($0.64\%/2 = 0.32\%$) switch to ownership under the First Time Homeownership Subsidy, while the “Lower Dispersion” series sets $\sigma_{\omega,B}$ so that twice as many renters ($2 \times 0.64\% = 1.28\%$) switch. For data definitions see notes for Figure 1 and Table 2.

Figure A.2: Credit Relaxation by Construction Supply Elasticity



Notes: Figure shows responses to our Credit Relaxation experiment (as also shown in Figure 7a) comparing the Benchmark model to a model with no construction sector and a fixed housing supply \bar{H} . For data definitions see notes for Figure 1 and Table 2. Additional definitions are "House Price" (p_t), "Rent" (q_t), " H " (H_t).

B Empirical Appendix

This section describes our data construction and presents additional empirical results and robustness checks.

B.1 Data Construction

We construct three different data sets for each of the three instruments.

B.1.1 LS Instrument Data Set

We create an annual panel of 413 CBSAs and metropolitan divisions (henceforth CBSAs) from 1976 to 2018, although not all data is available for every CBSA in every year. We do our best to use data for each metropolitan division separately, but in some cases where data is not available at the metropolitan division level we use data for its parent CBSA. This is necessary for the FHFA HPI for Boston, Chicago, Detroit, Seattle, and Washington, DC and for the housing stock for Dallas, LA, New York, San Francisco, Detroit, Miami, Philadelphia, and Chicago. Additionally, some of the data we use is at a higher frequency. We either take annual averages or choose quarter 2 for a given year, as appropriate (quarter 2 for HPI, annual averages for most other variables).

Our data sources are:

- House Prices: Our primary data source is CoreLogic's single family combined (detached and non-detached) price index, which they call tier 11. This data set is proprietary and not included in our replication package but can be purchased from CoreLogic. Our monthly data covers 402 CBSAs from 1976 to 2018. We use FHFA house price indices and Freddie Mac house price indices at the CBSA level as a supplementary data source.
- Rents: Our rent series is the CBRE Economic Advisers Torto-Wheaton index. In particular, we use their nominal rent index. This is available for 66 geographic areas that we map to CBSAs from 1980 to 2017. We are able to map to a quarterly panel for 53 CBSAs beginning in 1989 and 62 CBSAs beginning in 1994.
- Homeownership rates:
 - Our main homeownership measure is from the Census Housing Vacancy Survey. The Census produces homeownership rates at the CBSA level, however the CBSA definitions change over time. They use 1980 MSA definitions from

1986-1994, 1990 MSA definitions from 1995-2004, 2000 CBSA definitions from 2005-2014, and 2010 CBSA definitions from 2015-2017. We use a crosswalk to link these longitudinally. To deal with changing definitions, we use data on homeownership rates aggregated from the county level to each MSA/CBSA definition. If the difference between the homeowner rates using the two different CBSA definitions is more than 5% in either of the two closest censuses, we drop all log changes that go across the definition break point from our analysis. For instance, the 1990 MSA definitions include far fewer suburbs in the New York Metropolitan Area than the 2000 CBSA definition. As a result, in the Census data the homeownership rate is 37% in 2004 and 55% in 2005. Using the Census data, we see that using both the 2000 and 2010 census data (the two closest censuses to the 2004-2005 switch), the difference between the homeownership rates based on the two definitions is over 45%, so we drop any log changes in the homeownership rate that cross over the 2004-2005 redefinition. For robustness, we add analysis at the state level where geographic definitions do not change.

- Our secondary homeownership measure is from the ACS one-year estimates at the CBSA level, which is downloaded from the Census.
- Credit Data: Our credit data come from the Home Mortgage Disclosure Act micro-data, which we collapse to the CBSA level. Our main measure is the dollar volume of loans originated, but in the appendix we use robustness to the number of loans and the loan to income ratio. We also use the HMDA data together with the history of national conforming loan limits from the FHFA to create our instrument, which is the fraction of originations within 5% of the CLL.
- Housing Units: Our housing units data come from the Census.
- Housing supply elasticity: We use data from (Saiz) which we crosswalk from his MSA definitions to our CBSAs using principal cities.
- Employment and industry shares: We use the annual series of county-level employment from the QCEW and aggregate to the CBSA level to create a measure of log employment and employment shares for each NAICS two-digit industry. The QCEW suppresses observations where employment in a county-year-industry is small. To handle cases where a county barely slips below the suppression threshold for one year, we linearly interpolate employment when we have a few missing years. For

other cases, employment is small enough for a missing year that ignoring the issue does not matter once we aggregate to the CBSA level.

One data source that merits additional discussion is the CBRE Torto-Wheaton rent index. As mentioned in the main text, it measures the average change in rents for identical units in the same multi-family buildings. This has two advantages. First, it is a “repeat sales” methodology while most rent measures (e.g., the BLS) tend to be average or median rents. Second, it focuses on newly rented units, which is more appropriate for a price-rent ratio. In unreported results, we have compared the TW index with several other rent measures and have found two main results. First, the TW rent index is far more volatile than average or median rent series that do not use rents for newly-rented units. This makes sense: average rent series include contracts negotiated a long time ago and also include properties where a landlord has not passed rent increases through to a tenant in order to keep a good tenant and avoid paying the costs of finding a new tenant. Second, one may be concerned that the TW rent index is not representative because it only includes large, multi-family buildings. To assuage this concern, we obtained a single family rent index from a major data vendor. While we are not permitted to publish results with this data, we found that it was highly correlated with the TW rent index.

B.1.2 DK Instrument Data Set

For the DK instrument, our base data set is a data file provided to us by DK. The DK data set is at the county level, which we collapse to the CBSA level weighting by population. We then merge in the CoreLogic HPI data, TW rent data, and Census and ACS homeownership rate data as described above. The analysis uses data from 2001 to 2010 as with DK, and yields observations for 287 CBSAs.

B.1.3 MS Instrument Data Set

For the MS data, we begin with the Mian-Sufi NCL share in 2002 provided to us by MS for 259 CBSAs. We merge in homeownership data from the ACS as our main measure as it is available for more CBSAs and our instrument is underpowered for the smaller subsample with Census Vacancy Survey data. However, this limits us to a post-2005 sample for homeownership. Unreported results using contracted rents from the ACS to create a price-rent ratio are similar because ACS rents are very sticky and thus not ideal. We also merge in the HMDA data, CoreLogic and FHFA house price indices, and units as described above for the LS instrument. Our final data set includes 245 CBSAs with house price data from 1990 to 2017 and homeownership data from 2005 to 2017.

B.2 LS Instrument Details and Robustness

The Loutskina-Strahan instruments are the interaction of the change in the national conforming loan limit and the share of HMDA mortgage originations within 5% of the conforming loan limit in the prior year, as well as the interaction of this first instrument with the Saiz housing supply elasticity. We use the CLL for single-unit mortgages provided by FHFA. As mentioned in a footnote in the main text, starting in 2008 Congress allowed the CLL to rise by more in high-cost cities if their local house price index grew sufficiently quickly. 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 change in the national CLL regardless of the change in the local CLL in high-cost areas.

B.2.1 First Stage

Table B.1 shows the first stage for the Loutskina-Strahan instrument for cities for which we can calculate a price-rent ratio. For the dollar value of loans, the number of loans, and the loan to income ratio, we have a positive coefficient on the first instrument, the interaction of the fraction near the CLL times the change in the CLL, and a negative coefficient of the second instrument, the interaction of the first instrument with the Saiz housing supply elasticity. This makes sense: the effect of a change in the CLL on credit supply is stronger in places with more loans near the CLL in the prior period and this effect is stronger in more inelastic cities.

The joint first stage F -statistic for the two instruments is between 6 and 10. To make sure our results are not biased by weak instruments, we supplement our results with a LIML estimation below as a robustness check.

B.2.2 Additional Results and Robustness

In this section, we present additional results and robustness for the Loutskina-Strahan instrument.

Figure B.1 shows the impulse response of rents, which we describe in the main text. The response of rents has a similar shape to house prices but a smaller magnitude and is less statistically significant. However, the rent response is substantial, which one would not find using a measure of average or median rents.

Our first robustness check is to use a balanced sample for all four of our outcome variables: the price-rent ratio, homeownership rate, house prices, and rents. This contrasts to

Table B.1: Loutskina-Strahan First Stage

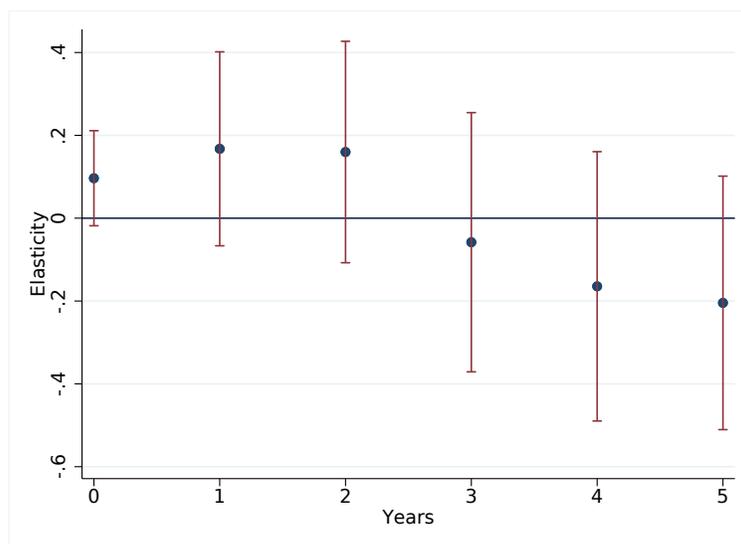
Dep Var	(1) $\Delta \log \left(\frac{\$ \text{ Volume}}{\text{of Loans}} \right)$	(2) $\Delta \log \left(\frac{\# \text{ of}}{\text{Loans}} \right)$	(3) $\Delta \log \left(\frac{\text{Loan}}{\text{Income}} \right)$	(4) $\log \left(\frac{\text{Price}}{\text{Rent}} \right)$
Frac Near CLL $\times \Delta CLL$ (Instrument 1)	43.802*** (12.04)	35.721** (11.72)	41.292*** (11.94)	3.771 (8.17)
Frac Near CLL $\times \Delta CLL \times \text{Saiz}$ (Instrument 2)	-62.037*** (15.85)	-56.109*** (15.89)	-52.607*** (14.88)	15.887 (11.36)
Frac Near CLL	-2.727** (0.99)	-2.863** (0.96)	-2.795** (1.08)	1.851*** (0.36)
Lag Frac Near CLL	2.797* (1.24)	3.704** (1.17)	2.950* (1.26)	-0.764* (0.34)
Lag Instrument 1	14.083 (10.60)	6.591 (11.37)	14.005 (11.62)	16.107 (8.92)
Lag Instrument 2	-22.405 (14.26)	-4.163 (15.21)	-17.983 (14.80)	-12.342 (12.90)
R^2	0.733	0.716	0.721	0.827
F On Instrument	7.733	6.295	6.640	9.583
N	1480	1480	1364	1480

Notes: This table shows the first stage of the Loutskina-Strahan instrument approach for CBSAs with non-missing price-rent ratios. The X variable is indicated in each column. Note that the first stage is the same for all horizons. In practice, we also include two lags of the outcome variable and two lags of the endogenous variable in the first stage; this turns out not to affect the results much so we omit these so we do not have to show a separate first stage for each outcome and endogenous variable. Standard errors are robust. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

the the main text, where we use all available data. As a result, we include some observations in the estimation of the homeownership rate, for instance, where we do not have data on price-rent ratios, and vice-versa. Limiting ourselves to CBSAs and periods where we have data on all of the outcomes gives the impulse responses shown in blue dots in Figure B.2 with 95% confidence intervals shown as red bars. The point estimates for original baseline specification in the main body of the paper are shown in orange circles. One can see that balancing the sample does not change our results significantly. We omit results for the inverse ratio as that specification by construction only includes observations where we observe the homeownership rate, rents, and prices, and is thus unaffected by this balancing.

Figures B.3 and B.4 show results when we use the number of loans and the loan-to-income ratio as the credit measure rather than the dollar value of loans, following Favara and Imbs (2015). Things look broadly similar in shape but not in magnitude. This makes sense: the reduced form between the instrument and the outcome remains the same, but the first stage by which it is scaled changes as we change the X variable.

Figure B.1: Loutskina-Strahan Instrument LP-IV Impulse Response For Rent



Notes: 95 % confidence interval shown in red bars. The figure shows panel local projection instrumental variables estimates of the response of log rents to dollar credit volume. The second stage is as indicated in equation (1) and the first stage is as indicated in equation (2). The two instruments are $ShareNearCLL_{i,t} \times \%ChangeInCLL_t$. and $ShareNearCLL_{i,t} \times \%ChangeInCLL_t \times Z(SaizElasticity_i)$. Control variables include $ShareNearCLL_{i,t}$ and its lag, lags of both instruments, and two lags of both the outcome variable and the endogenous variable. All standard errors are clustered by CBSA.

Next, we address a potential issue with our data: the homeownership rate measure from the Census changes MSA/CBSA definitions every ten years. In most cases, this is not a problem because the CBSA definitions do not change substantially. For example, a redefinition might imply that only a small, rural county is added as the MSA/CBSA grows, leaving the homeownership rate largely unchanged. However, in a few cases the CBSA definition changes lead to abrupt jumps in the homeownership rate. A good example is New York, where a definition change that added many suburbs dramatically raised the homeownership rate. Including these periods in our estimates would likely lead to significant noisiness and instability in our estimates.

Our approach to dealing with these issues in the main text is to check whether the change in definitions alters the homeownership rate by more than 5% in adjacent decennial censuses. If it does, we treat the area as a different CBSA before and after the redefinition. If not, we leave it as a single CBSA through the redefinition. This approach has the advantage of nonparametrically allowing for jumps at the time or redefinition while still preserving the full sample of data.

In Figure B.5, we take an alternate approach and drop all CBSAs that have a jump of

more than 5% due to definition changes at any point in the panel. This reduces the number of CBSAs – and particularly large CBSAs like New York, LA, Chicago, and Miami which are more likely to have definitional changes – but yields a more consistent homeownership rate series for the CBSAs that remain. Figure B.5 shows the results from this alternative specification, with the baseline results overlaid in orange. We observe that the impulse responses of the price to rent ratio, homeownership rates, and house prices are modestly stronger, while the inverse supply slope regression preserves our finding of a steep tenure supply slope, albeit with a somewhat different time pattern.

In Figure B.6, we add log employment and employment shares for all 2-digit NAICS industries from the QCEW as controls. These time-varying measures account for changes in the industrial structure and employment that are not picked up by the CBSA fixed effects. One can see that the results change little from our baseline specification.

Figure B.7 shows estimates using LIML rather than two-stage least squares. The point estimates are very similar, with the exception of the house price response at the 3-year horizon where both the LIML point estimate and confidence interval blow up. Overall, the similarity with two-stage least squares suggests that weak instruments are not biasing our results.

Finally, Figures B.8 and B.9 show results using slightly different data than our main analysis. Figure B.8 uses the FHFA purchase-only house price index rather than the CoreLogic house price index. One can see a similarly shaped IRF, although the magnitudes are larger with the FHFA index. As a result, the inverse supply slope is similar, although slightly smaller. Figure B.9 uses ACS homeownership rate data rather than Census. The ACS has the advantage of having consistent geographies but it only starts in 2005, so we lose a lot of the variation in the instruments in the boom (there is very little variation in the bust since the conforming loan limit only adjusts upwards and was thus flat throughout the bust). The standard errors are thus quite wide, especially at the 2-year horizon, but the point estimates are generally of a similar magnitude. The inverse slope impulse response is similar although slightly smaller in magnitude until year 3, after which the estimates are imprecise. Overall, our results are robust to the data series used.

B.3 DK Instrument Details and Robustness

As described in the main body, we use a county-level data set generously provided by Di Maggio and Kermani which we collapse to the CBSA level. We are able to run the same regression at the CBSA level that they run at the county level with the exception of one proprietary control variable: the share of loans that are subprime. We also use log

changes rather than percent changes. We run the same regression as DK, then transform the impulse response from log changes to log levels by cumulating the coefficients from the log changes regression.

B.3.1 Replicating DK's Exact Specification

Figure B.10 replicates DK's exact specification by showing the impulse response in growth rates rather than levels. One can see that house price growth peaks after two years (in 2005) and then mean reverts slightly (although not statistically significantly). DK find a longer, more gradual boom and no mean reversion. The difference is essentially entirely due to collapsing to the CBSA level and dropping non-CBSA counties.

B.3.2 Robustness

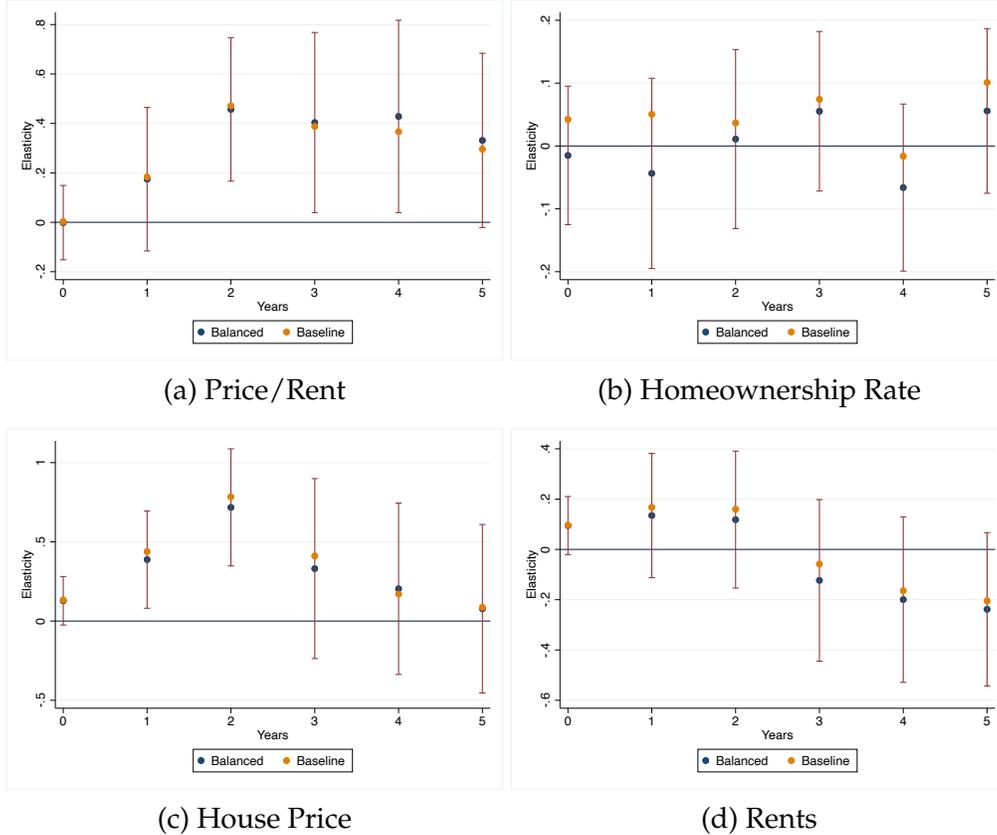
Figures B.11 and B.12 show two robustness checks for the DK approach. The first drops all controls except for fixed effects; the results are quite similar. The second controls for log employment and 2-digit industry shares. This causes the standard errors to widen considerably. On the whole, though, the results are robust to these changes in controls.

B.4 MS Instrument Details and Robustness

As mentioned in the main text, while Mian and Sufi use 2002 as a base year, we are unable to do so because the ACS only starts in 2005. However, for our house prices, we do not face this limitation, and could have used a 2002 base. Figure B.13 shows the results from this alternative regression with a 2002 base year. One can see that the impulse response for house prices returns to zero in 2013, motivating our use of 2013 as the base year in the main text.

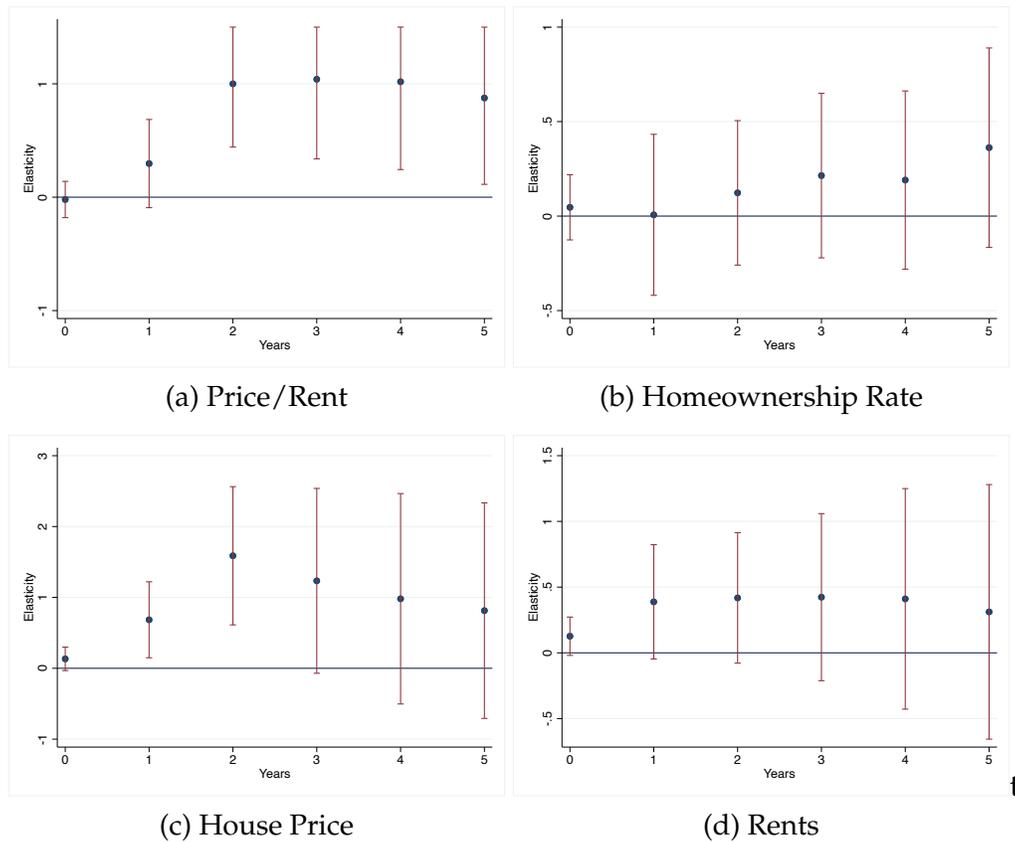
Figure B.14 reestimates our main MS regression with log employment and industry share controls. We find responses with a similar shape to the baseline, but with slightly lower point estimates, and much wider standard errors, particularly for house prices. The effect on house prices peaks at 0.78 while the effect on the homeownership rate peaks at 0.20, implying a naive slope of 4, consistent with the range of slope estimates in our earlier findings. Overall, we find that our MS results are robust to these controls.

Figure B.2: Loutskina-Strahan Instrument LP-IV Impulse Responses: Balanced Sample



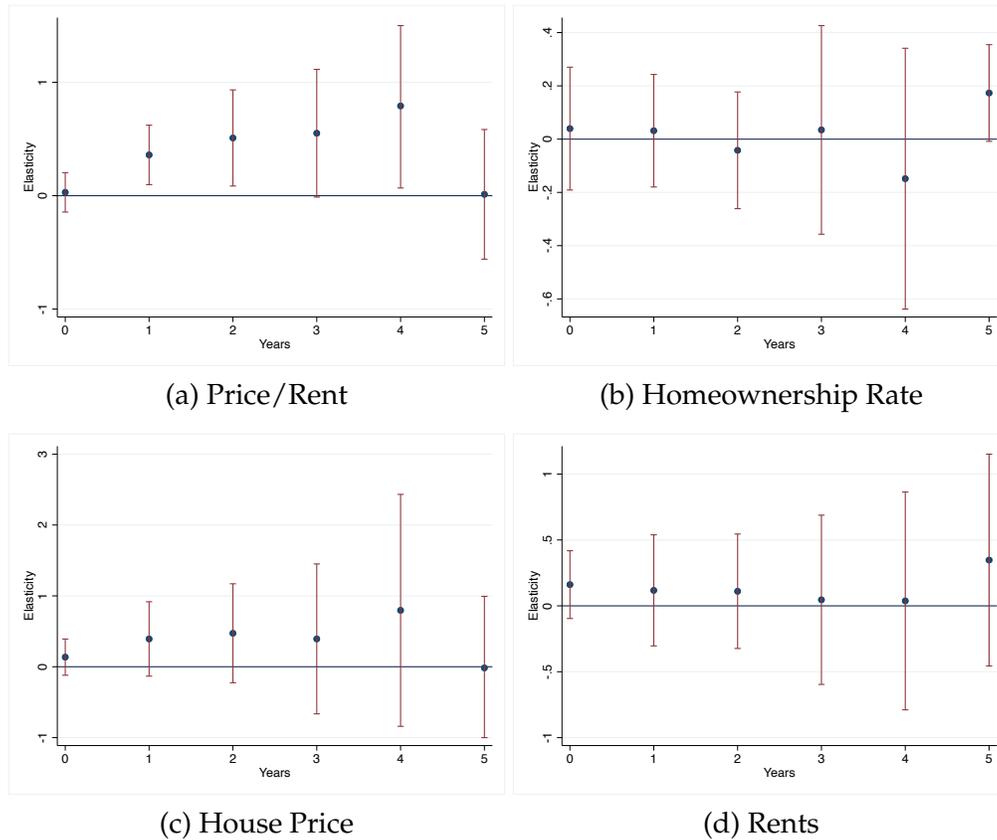
Notes: Blue dots indicate the balanced samples across all outcomes and 95 % confidence interval shown in red bars. The original baseline point estimates in the main body of the paper are shown in orange dots. The figure shows panel local projection instrumental variables estimates of the response of the indicated outcomes to dollar credit volume. For panels A to C, the second stage is as indicated in equation (1) and the first stage is as indicated in equation (2). The two instruments are $ShareNearCLL_{i,t} \times \%ChangeInCLL_t$. and $ShareNearCLL_{i,t} \times \%ChangeInCLL_t \times Z(SaizElasticity_i)$. Control variables include $ShareNearCLL_{i,t}$ and its lag, lags of both instruments, and two lags of both the outcome variable and the endogenous variable. For panel D, we use the same controls but the outcome variable is the homeownership rate and the credit variable is replaced with log price-rent ratio at time $t + k$ to obtain a coefficient for the inverse supply curve slope. All standard errors are clustered by CBSA.

Figure B.3: Loutskina-Strahan Instrument LP-IV Impulse Responses: Number of Loans as Outcome



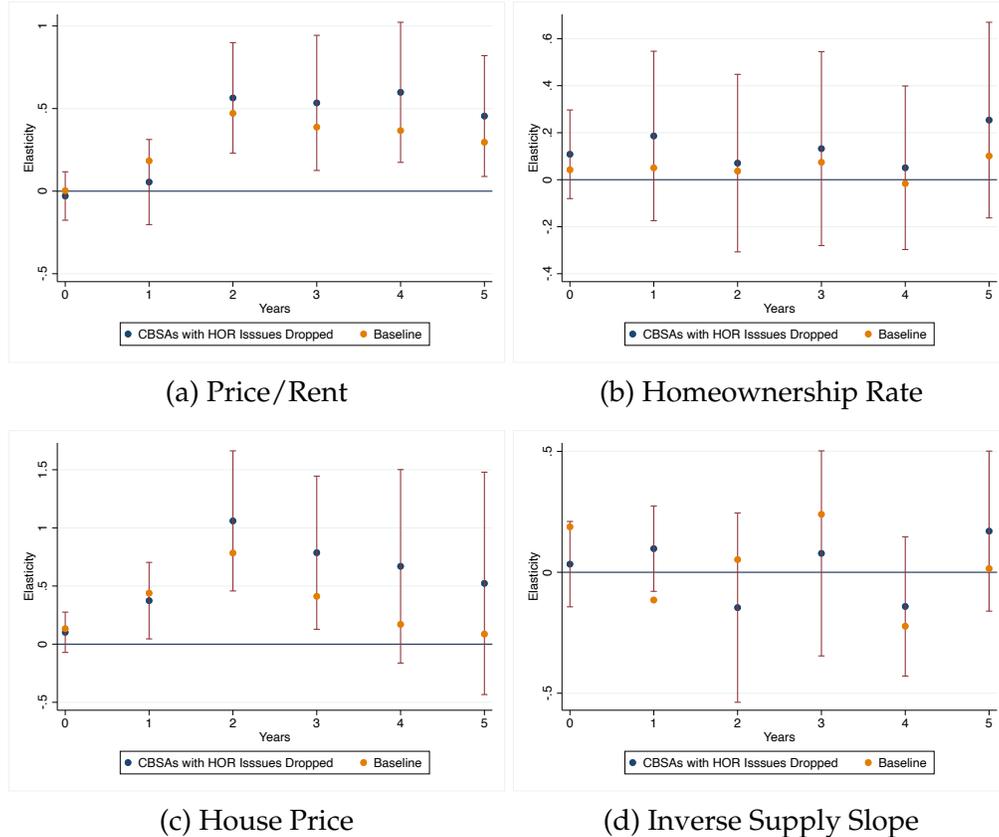
Notes: 95 % confidence interval shown in red bars. The figure shows panel local projection instrumental variables estimates of the response of the indicated outcomes to the number of new loans. For panels A to C, the second stage is as indicated in equation (1) and the first stage is as indicated in equation (2). The two instruments are $ShareNearCLL_{i,t} \times \%ChangeInCLL_t$ and $ShareNearCLL_{i,t} \times \%ChangeInCLL_t \times Z(SaizElasticity_i)$. Control variables include $ShareNearCLL_{i,t}$ and its lag, lags of both instruments, and two lags of both the outcome variable and the endogenous variable. For panel D, we use the same controls but the outcome variable is the homeownership rate and the credit variable is replaced with log price-rent ratio at time $t + k$ to obtain a coefficient for the inverse supply curve slope. All standard errors are clustered by CBSA.

Figure B.4: Loutskina-Strahan Instrument LP-IV Impulse Responses: Loan to Income Ratio as Outcome



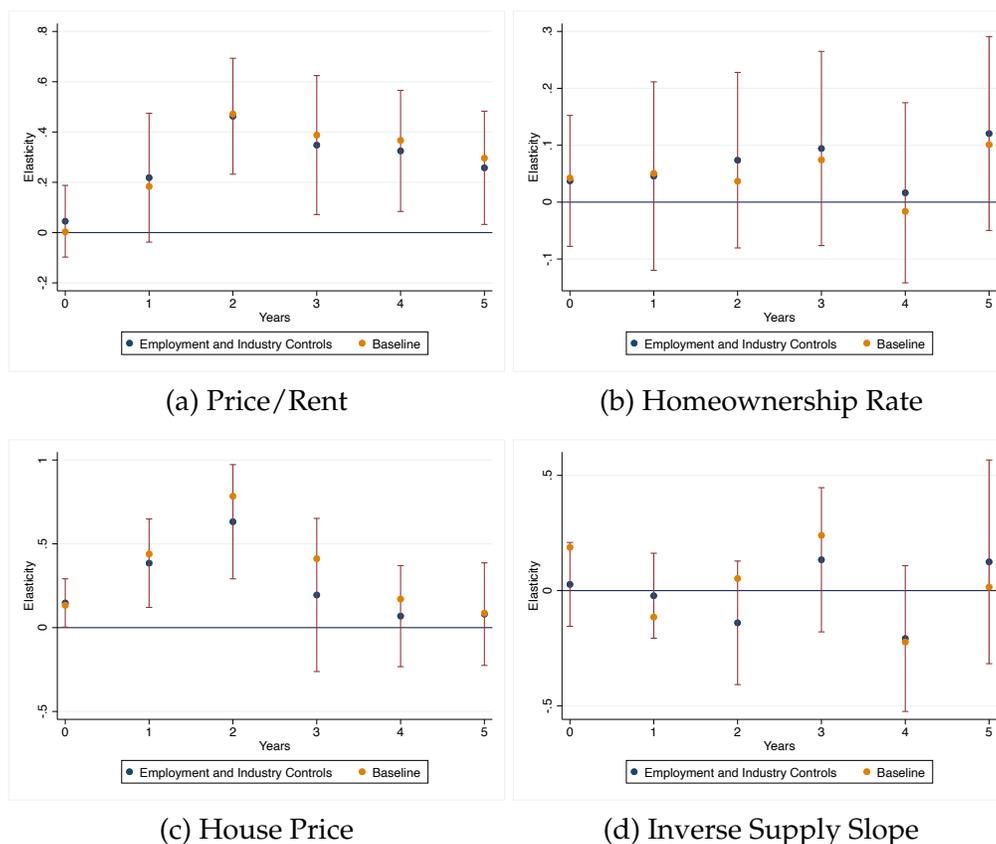
Notes: 95 % confidence interval shown in red bars. The figure shows panel local projection instrumental variables estimates of the response of the indicated outcomes to the loan to income ratio with loans coming from HMDA and income coming from the IRS. For panels A to C, the second stage is as indicated in equation (1) and the first stage is as indicated in equation (2). The two instruments are $ShareNearCLL_{i,t} \times \%ChangeInCLL_t$ and $ShareNearCLL_{i,t} \times \%ChangeInCLL_t \times Z(SaizElasticity_i)$. Control variables include $ShareNearCLL_{i,t}$ and its lag, lags of both instruments, and two lags of both the outcome variable and the endogenous variable. For panel D, we use the same controls but the outcome variable is the homeownership rate and the credit variable is replaced with log price-rent ratio at time $t + k$ to obtain a coefficient for the inverse supply curve slope. All standard errors are clustered by CBSA.

Figure B.5: Loutskina-Strahan Instrument LP-IV Impulse Responses: Dropping All CBSAs With Issues With Homeownership Rate Instead of Dropping Just Bad Periods



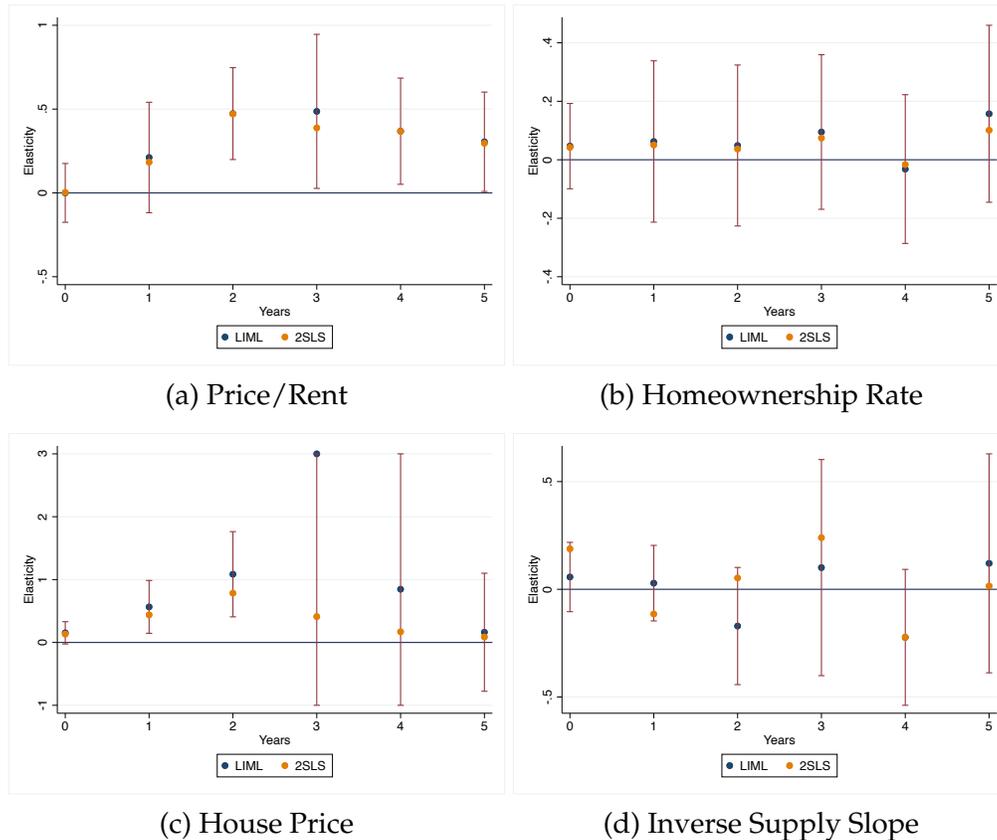
Notes: Blue dots indicate the the estimates dropping CBSAs with changes in the homeownership rate due to changes in the Census definition of the CBSA and 95 % confidence interval shown in red bars. The original baseline point estimates in the main body of the paper are shown in orange dots. The figure shows panel local projection instrumental variables estimates of the response of the indicated outcomes to dollar credit volume. For panels A to C, the second stage is as indicated in equation (1) and the first stage is as indicated in equation (2). The two instruments are $ShareNearCLL_{i,t} \times \%ChangeInCLL_t$. and $ShareNearCLL_{i,t} \times \%ChangeInCLL_t \times Z(SaizElasticity_i)$. Control variables include $ShareNearCLL_{i,t}$ and its lag, lags of both instruments, and two lags of both the outcome variable and the endogenous variable. For panel D, we use the same controls but the outcome variable is the homeownership rate and the credit variable is replaced with log price-rent ratio at time $t + k$ to obtain a coefficient for the inverse supply curve slope. All standard errors are clustered by CBSA.

Figure B.6: Loutskina-Strahan Instrument LP-IV Impulse Responses: Log(Employment) and Two-Digit Industry Shares Controls



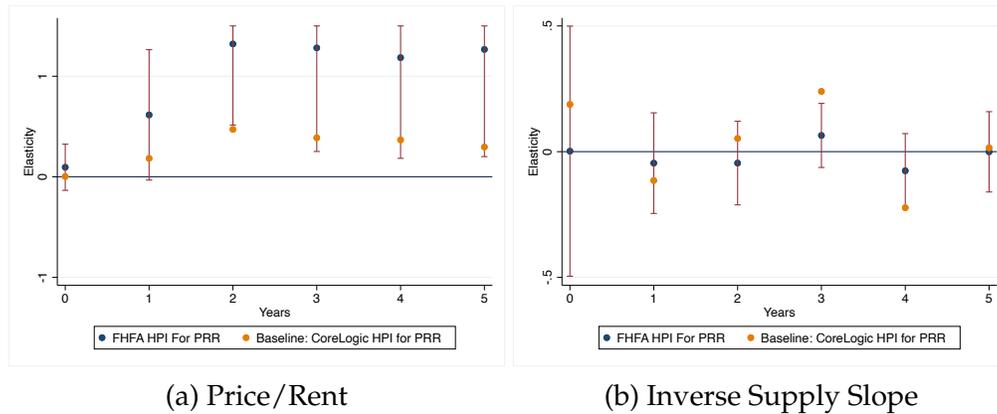
Notes: Blue dots indicate estimates with log employment and industry share controls from QCEW and 95 % confidence interval shown in red bars. The original baseline point estimates in the main body of the paper are shown in orange dots. The figure shows panel local projection instrumental variables estimates of the response of the indicated outcomes to dollar credit volume. For panels A to C, the second stage is as indicated in equation (1) and the first stage is as indicated in equation (2). The two instruments are $ShareNearCLL_{i,t} \times \%ChangeInCLL_t$ and $ShareNearCLL_{i,t} \times \%ChangeInCLL_t \times Z(SaizElasticity_i)$. Control variables include $ShareNearCLL_{i,t}$ and its lag, lags of both instruments, and two lags of both the outcome variable and the endogenous variable. For panel D, we use the same controls but the outcome variable is the homeownership rate and the credit variable is replaced with log price-rent ratio at time $t + k$ to obtain a coefficient for the inverse supply curve slope. All standard errors are clustered by CBSA.

Figure B.7: Loutskina-Strahan Instrument LP-IV Impulse Responses: LIML



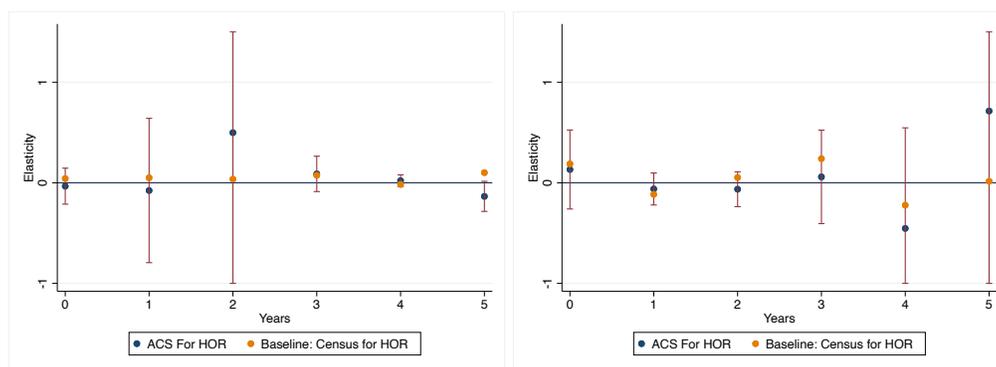
Notes: Blue dots indicate LIML and 95 % confidence interval shown in red bars. The original baseline point estimates in the main body of the paper estimated by two-stage least squares are shown in orange dots. The figure shows panel local projection instrumental variables estimates of the response of the indicated outcomes to dollar credit volume. For panels A to C, the second stage is as indicated in equation (1) and the first stage is as indicated in equation (2). The two instruments are $ShareNearCLL_{i,t} \times \%ChangeInCLL_t$, and $ShareNearCLL_{i,t} \times \%ChangeInCLL_t \times Z(SaizElasticity_i)$. Control variables include $ShareNearCLL_{i,t}$ and its lag, lags of both instruments, and two lags of both the outcome variable and the endogenous variable. For panel D, we use the same controls but the outcome variable is the homeownership rate and the credit variable is replaced with log price-rent ratio at time $t + k$ to obtain a coefficient for the inverse supply curve slope. All standard errors are clustered by CBSA.

Figure B.8: Loutskina-Strahan Instrument: FHFA House Price Index



Notes: Blue dots indicate price-rent ratios made with FHFA purchase only house price index and 95 % confidence interval shown in red bars. The original baseline point estimates in the main body of the paper which use the CoreLogic house price index are shown in orange dots. The figure shows panel local projection instrumental variables estimates of the response of the indicated outcomes to dollar credit volume. For panel A, the second stage is as indicated in equation (1) and the first stage is as indicated in equation (2). The two instruments are $ShareNearCLL_{i,t} \times \%ChangeInCLL_t$ and $ShareNearCLL_{i,t} \times \%ChangeInCLL_t \times Z(SaizElasticity_i)$. Control variables include $ShareNearCLL_{i,t}$ and its lag, lags of both instruments, and two lags of both the outcome variable and the endogenous variable. For panel B, we use the same controls but the outcome variable is the homeownership rate and the credit variable is replaced with log price-rent ratio at time $t + k$ to obtain a coefficient for the inverse supply curve slope. All standard errors are clustered by CBSA.

Figure B.9: Loutskina-Strahan Instrument: ACS Homeownership Rate

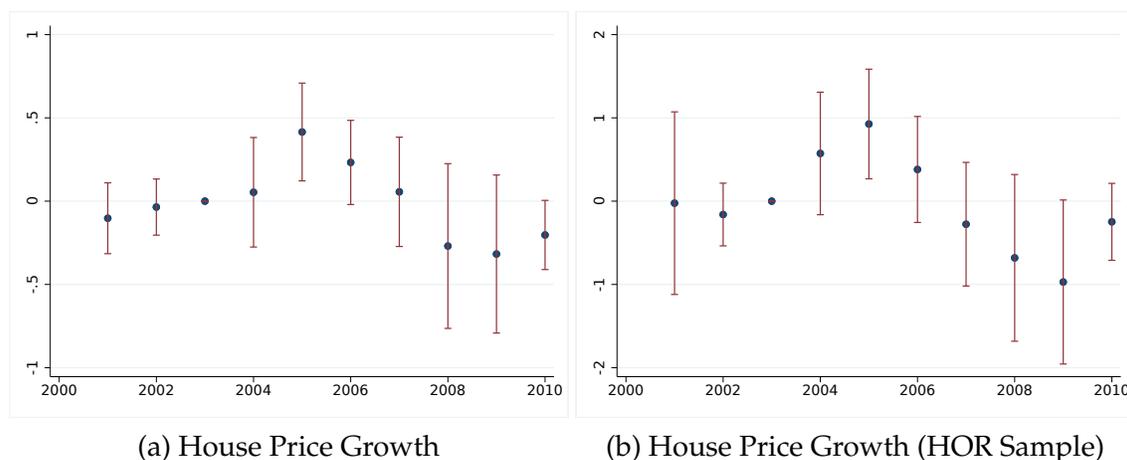


(a) Homeownership Rate

(b) Inverse Supply Slope

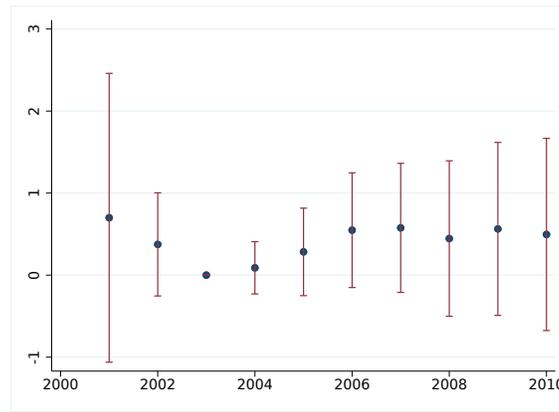
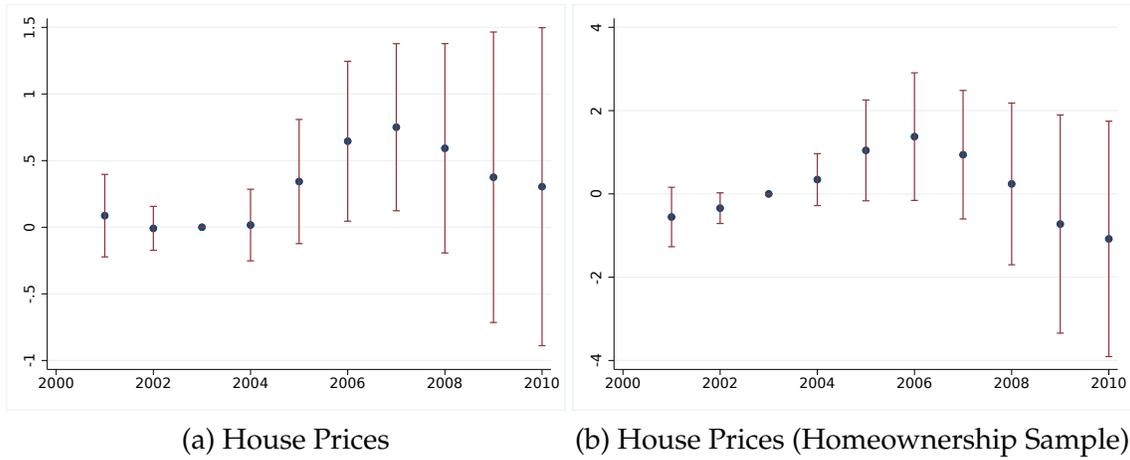
Notes: Blue dots indicate homeownership rates made with ACS (starting in 2005) with 95 % confidence interval shown in red bars. The original baseline point estimates in the main body of the paper which use Census homeownership rates are shown in orange dots. The figure shows panel local projection instrumental variables estimates of the response of the indicated outcomes to dollar credit volume. For panel A , the second stage is as indicated in equation (1) and the first stage is as indicated in equation (2). The two instruments are $ShareNearCLL_{i,t} \times \%ChangeInCLL_t$. and $ShareNearCLL_{i,t} \times \%ChangeInCLL_t \times Z(SaizElasticity_i)$. Control variables include $ShareNearCLL_{i,t}$ and its lag, lags of both instruments, and two lags of both the outcome variable and the endogenous variable. For panel B, we use the same controls but the outcome variable is the homeownership rate and the credit variable is replaced with log price-rent ratio at time $t + k$ to obtain a coefficient for the inverse supply curve slope. All standard errors are clustered by CBSA.

Figure B.10: Di Maggio-Kermani: IRFs in Growth Rates as in DK Specification



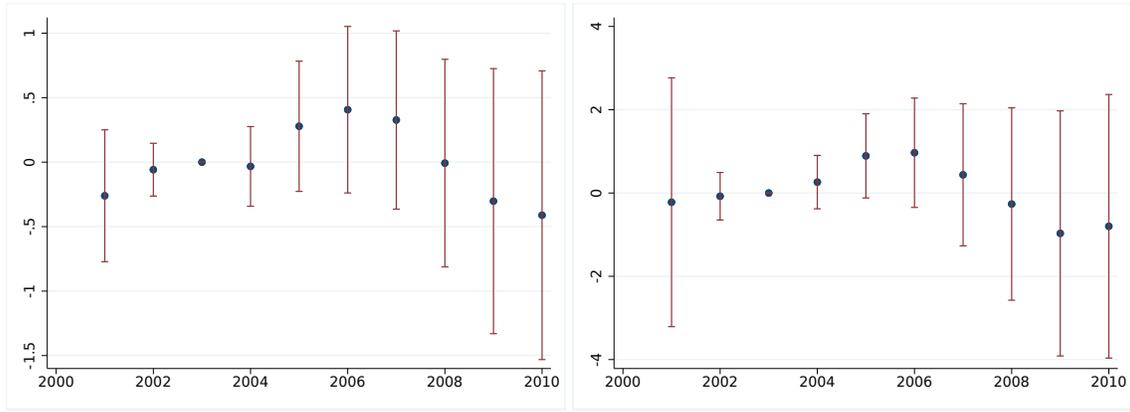
Notes: 95 % confidence interval shown in red bars. Figures show estimates of β_k for each indicated year estimated from equation (3), with the instrument being $Z_i = APL_{2004} \times OCC_{2003}$ and 2003 being the base year. The regression is weighted by population and standard errors are clustered by CBSA. The controls include median income growth, population growth, the Saiz (2010) elasticity interacted with a dummy for post-2004, the fraction of loans originated by HUD-regulated lenders interacted with a dummy for post-2004, and the fraction of HUD-regulated lenders interacted with a dummy for APLs. The controls are as Di Maggio and Kermani (2017) except our data is at the CBSA level and omits a control for the fraction of loans originated that are subprime (FICO under 620), which is based on proprietary data. All regressions are weighted by population and standard errors are clustered by CBSA as in the DK paper.

Figure B.11: Di Maggio-Kermani: Without Controls



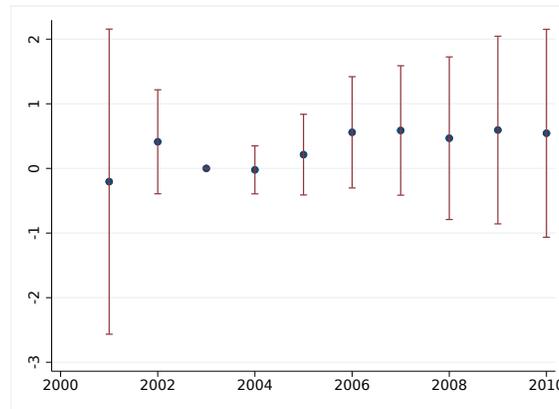
Notes: 95 % confidence interval shown in red bars. Panels A to C show estimates of the cumulative sum from 2003 of β_k for each indicated year estimated from equation (3), with the instrument being $Z_i = APL_{2004} \times OCC_{2003}$ and 2003 being the base year. The regression is weighted by population and standard errors are clustered by CBSA. The regression does not include the DK controls. All regressions are weighted by population and standard errors are clustered by CBSA as in the DK paper.

Figure B.12: Di Maggio-Kermani:With Log Employment and Industry Share Controls



(a) House Prices

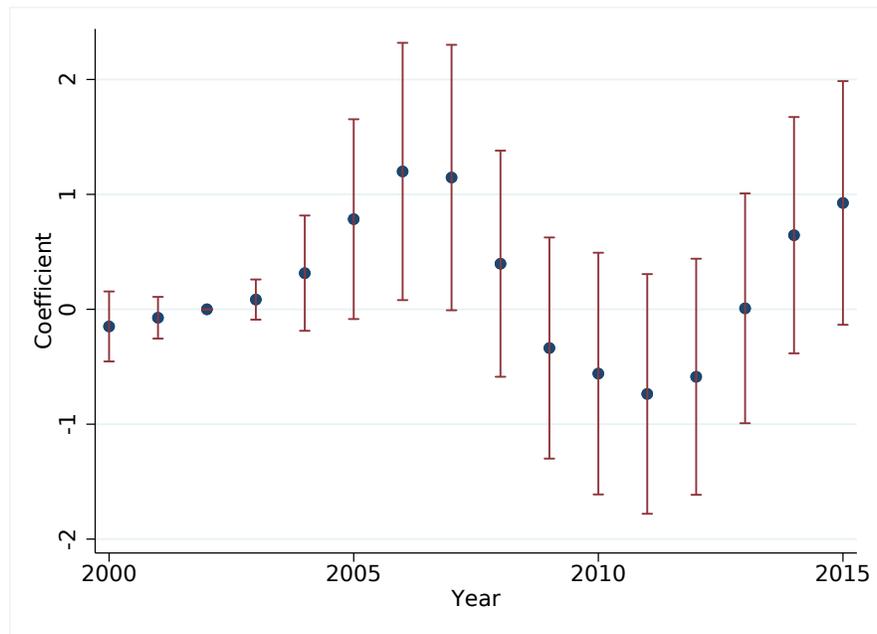
(b) House Prices (Homeownership Sample)



(c) Homeownership Rate

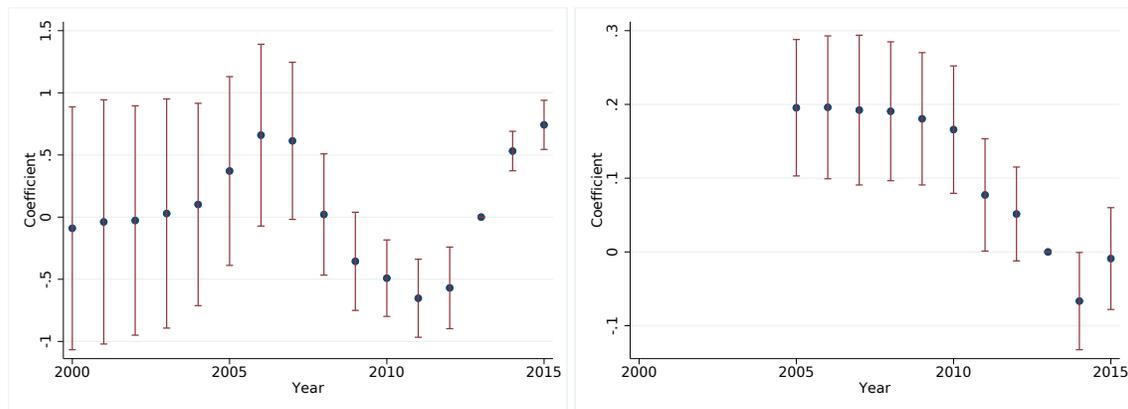
Notes: 95 % confidence interval shown in red bars. Panels A to C show estimates of the cumulative sum from 2003 of β_k for each indicated year estimated from equation (3), with the instrument being $Z_i = APL_{2004} \times OCC_{2003}$ and 2003 being the base year. The regression is weighted by population and standard errors are clustered by CBSA. The controls include median income growth, population growth, the Saiz (2010) elasticity interacted with a dummy for post-2004, the fraction of loans originated by HUD-regulated lenders interacted with a dummy for post-2004, and the fraction of HUD-regulated lenders interacted with a dummy for APLs. These controls are as Di Maggio and Kermani (2017) except our data is at the CBSA level and omits a control for the fraction of loans originated that are subprime (FICO under 620), which is based on proprietary data. We also include log employment and 2-digit industry shares from the QCEW as controls in this figure. All regressions are weighted by population and standard errors are clustered by CBSA as in the DK paper.

Figure B.13: Mian-Sufi PLS Expansion Reduced Form With 2002 Base For House Prices



Notes: 95 % confidence interval shown in red bars. Figure shows estimates of the effect of a city's NCL share on each outcome based on estimating equation 3 with the instrument being $Z_i = NCLShare_i^{2002}$ and 2003 being the base year. All standard errors are clustered by CBSA as in the MS paper.

Figure B.14: Mian-Sufi PLS Expansion Reduced Form With Log Employment and Industry Share Controls



(a) House Prices

(b) Homeownership Rate

Notes: 95 % confidence interval shown in red bars. Each panel shows estimates of the effect of a city's NCL share on each outcome based on estimating equation 3 with the instrument being $Z_i = NCLShare_i^{2002}$ and 2013 being the base year. All panels include log employment and 2-digit industry shares from the QCEW as controls. All standard errors are clustered by CBSA as in the MS paper.

B.5 National Price-Rent Ratio Construction

Our measure of the national price-rent ratio comes from the BEA and Flow of Funds. The ideal measure would be the ratio of the market value of household real estate (FRED code: BOGZ1FL155035013Q) to the value of owner-occupied housing services (FRED code: A2013C1A027NBEA). However, this housing service measure is only available annually. To obtain a quarterly measure, we use the fact that total housing services (FRED code: DHUTRC1Q027SBEA) are available quarterly and can serve as a proxy. With these series in hand, our construction proceeds as follows:

$$\begin{aligned} OwnerServices_t &= \left(\frac{OwnerServices_t}{TotalServices_t} \right) TotalServices_t \\ &= \left(\frac{OwnerServicesPerUnit_t}{TotalServicesPerUnit_t} \right) \left(\frac{OwnerUnits_t}{TotalUnits_t} \right) TotalServices_t. \end{aligned}$$

The first term in this expression is unknown at quarterly frequency, and will need to be approximated. The second term is the homeownership rate, which is available in the data quarterly, and the third term is total housing services, which is also available in the data quarterly. As a result, to obtain our measure of the homeownership rate, we interpolate the ratio $(OwnerServicesPerUnit_t / TotalServicesPerUnit_t)$ in logs between annual observations. We then compute our quarterly owner services series as

$$OwnerServices_t = \left(\frac{\widehat{OwnerServicesPerUnit_t}}{TotalServicesPerUnit_t} \right) \left(\frac{OwnerUnits_t}{TotalUnits_t} \right) TotalServices_t.$$

where the hat denotes the interpolated value of this ratio.