

Do Credit Conditions Move House Prices?

Daniel L. Greenwald* and Adam Guren[‡]

February 2024

Abstract

Did credit drive the 2000s housing cycle? The existing literature's findings range from credit having no effect to credit explaining most of the cycle. We show that these disparate results hinge on the extent to which landlords absorb credit-driven demand, which depends on the degree of housing market segmentation. We develop a model that nests cases between the extremes of no segmentation and perfect segmentation typically considered, estimate an elasticity that pins down the degree of segmentation, and use it to calibrate our model. We find that credit standards played an important role, explaining 32% to 53% of the boom.

*NYU Stern School of Business, dlg340@stern.nyu.edu.

[†]Boston University and NBER, guren@bu.edu.

[‡]We would like to thank Joe Vavra, Monika Piazzesi, Tim McQuade, Jaromir Nosal, Nelson Lind, David Berger, Mikhail Golosov (editor), three anonymous referees, 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, VMACS, Yale, USC, Northwestern, Columbia Business School, and the Federal Reserve Bank of Minneapolis 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. Adamson Bryant, Paul Dai, Liang Zhong, and Fanwen Zhu provided excellent research assistance. Guren thanks the National Science Foundation (grant SES-1623801) for financial support.

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 studies, 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. Others, 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

In this paper we attempt to reconcile these divergent findings to quantitatively assess the role of credit in driving the 2000s housing cycle.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, 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 and price-rent ratios. 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, fixing house prices to this level. Since credit conditions have a small effect on present and future rents, credit has a minimal impact on house prices in these models. Instead, when credit ex-

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 (2022), 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.

pands, households buy housing from landlords, increasing the homeownership rate.

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 levels of segmentation 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.

Measuring the slope of the tenure supply curve is complicated by simultaneous supply and demand shocks and measurement error. We follow the traditional approach of using demand instruments, namely instruments for credit supply to home purchasers.³ 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. This approach builds on an empirical literature that uses exogenous changes in credit to show that credit does affect house prices; our contribution is to provide a mapping from these local average treatment effects to calibrated models that allow researchers to quantitatively assess the aggregate role of credit in the housing cycle.

To this end, we employ three identification strategies from the existing literature. 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 (2022), 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.

In pursuing these three identification strategies, we use two measures of the local homeownership rate. First, we use the Census Housing and Vacancy Survey (HVS). While publicly available, the HVS is limited in its geographic coverage, suffers from changing geographies, and appears noisy. To address this, we introduce a new local homeownership rate constructed at the individual property level from micro data on the owners and residents of most housing units in the United States. Our measure has wider geographic coverage and is less noisy, which vastly improves the statistical precision of our estimates.

Despite relying on different sources of identification and operating through differ-

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

ent segments of the mortgage market, our three empirical approaches deliver 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 effects on the homeownership rate. Our preferred specifications deliver estimates of the slope that are at least 3.8 and are in most cases substantially larger, implying a high 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 heterogeneity 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 and then use the model to determine the role of credit in driving the 2000s housing boom.

We find that a relaxation of credit standards in isolation explains 32% of the rise in the price-rent ratio observed in the boom, with a lower bound of 22% accounting for parameter uncertainty. These results contrast to -2% explained by the model with no segmentation, and the 36% 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. Adding a 2pp decline in mortgage rates alongside the relaxation of credit standards allows our benchmark model to explain 70% of the observed rise in price-rent ratios, compared to 2% under no segmentation and 77% under full segmentation.

Our findings also have important implications for macroprudential policies that restrict credit standards, which have been adopted by many countries to rein in housing booms. These policies can be successful in slowing house price growth, but only in the presence of strong segmentation as we find. To show this quantitatively, we augment our main credit expansion experiment with additional supply and demand shocks reflecting factors such as expectations so that the model explains the entire boom in price-rent ratios and homeownership. From this starting point, we consider a counterfactual in which we prevent credit standards from relaxing during the boom, which reduces the observed rise in price-rent ratios by 53%. In contrast, the same policy exercise in a model with no segmentation would reduce the rise in price-rent ratios by just 3%.

We conclude by relaxing two important assumptions in our baseline model. First, we extend the model to allow landlords to use credit. While this amplifies the response of house prices to credit because the tenure supply curve now shifts outward, we find this

effect is quantitatively small after we recalibrate our model to match our key empirical moment. Second, we relax the segmentation between borrowers and savers, allowing them to freely trade housing. This dampens the response of house prices to credit shocks as savers play a role similar to the landlords in our main analysis. Nonetheless, changes to credit standards and interest rates still explain 51% of the observed rise in price-rent ratios in the boom, relative to 70% in the baseline model. Since this extension ignores real-world frictions due to indivisibility and variation in the quality and location of housing by allowing for highly unrealistic intensive margin adjustments by savers, we consider this an extremely conservative estimate of the effect of credit on house prices.

Related Literature. Empirically, our paper 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 (2023), Favara and Imbs (2015), Loutskina and Strahan (2015), Di Maggio and Kermani (2017), and Mian and Sufi (2022). Importantly, it is not straightforward to map the specific and typically cross-sectional quasi-random variation these papers exploit into the quantitative effect of credit on house prices in response to a particular set of structural shocks. We show how to combine estimated responses of the price-rent ratio and the homeownership rate to calibrate a structural model, which can then be used to assess the general equilibrium response of house prices to the shocks typically considered in analyses of the 2000s cycle. The closest counterpart to our empirical approach is Gete and Reher (2018), who also measure the impact of an identified credit shock on both the price-rent ratio and homeownership. They estimate a response of the price-rent ratio that is 85 times larger than the response of the homeownership rate, in line with our findings of strong segmentation.

Our work also relates to a literature using quantitative models to study the effect of credit supply on house prices, including 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 (2022), Justiniano et al. (2019), Liu et al. (2019), and Huo and Rios-Rull (2016). We explain the wide variation in results in this literature and provide a new empirical framework to reconcile them.

Our paper is particularly close to Landvoigt et al. (2015). We assume full segmentation between savers and borrowers in our baseline model (which we relax in extensions) and show how frictions between landlords and borrowers are critical for the effect of credit on house prices. Landvoigt et al. (2015) take the reverse approach, implicitly assuming full segmentation between borrowers and landlords by using a model without a tenure choice, and using a rich assignment model to show how sorting of owners with hetero-

geneous wealth across houses of different quality creates endogenous segmentation between borrowers and savers. We see these works as highly complementary, as both types of segmentation appear necessary for credit to have a strong effect on house prices.

Last, our modeling framework 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 Rognlie, Shleifer and Simsek (2018). Our contribution is to provide a tractable methodology for incorporating fractional and time-varying homeownership rates and a new empirical moment to discipline model calibration.

1 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.

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 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. This pattern changed dramatically during the 1998-2006 housing boom, during which the price-rent ratio and homeownership rate increased in tandem. Following the start of the bust in 2007, these variables reverse course, traveling a downward path and ending close to the typical values from the pre-boom era.

To understand the forces behind these patterns, we present a simple supply and demand treatment that illustrates the key mechanism at work in our model. As in Figure 1, we use the price-rent ratio on the y-axis and the homeownership rate on the x-axis, which represent the relative price and relative quantity of owned vs. rented housing. This focus on relative rather than absolute prices and quantities of housing ensures that changes are driven by movements in share of housing that is owner-occupied rather than from changes in the absolute quantity of housing units via construction.

The *tenure demand* curve represents the price schedule of households for owned vs.

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).


rented housing.⁴ As the price-rent ratio rises, more households prefer renting to owning, creating a downward slope. The *tenure supply* curve is the price schedule at which landlords are willing to sell rental housing to households as owner-occupied housing. The slope of the tenure supply curve reflects the marginal willingness of landlords to sell more units as the price-rent ratio rises, while shifts in the tenure supply curve reflect shocks that change landlords’ fundamental values of houses relative to rents.

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 nests models with no rental housing (e.g., Favilukis et al. (2017), Greenwald (2018), and Justiniano, Primiceri and Tambalotti (2015)) as well as models in which owning and renting are completely segmented, implying a fixed homeownership rate. 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 for now

⁴We follow the Census in using the term “tenure” to reflect whether housing is owned or rented.

Figure 2: Supply and Demand Intuition



is assumed to only affect households. This shifts tenure demand outward from the solid curve D to the dashed curve D' , as improvements in the cost or availability of financing induce more households to purchase rather than rent at a given relative price. The expansion of tenure demand due to the relaxation of credit produces a large rise in house prices because households in this model have no one to trade with except each other. However, this model cannot produce the rise in homeownership observed in Figure 1.

Panel (b) considers the alternate 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 tenure demand increases the homeownership rate but not the price-rent ratio. Reproducing the joint empirical pat-

tern 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 expected present value of rents in this model, a shock to future rents, such as the shock to future rental beliefs used in Kaplan et al. (2020), can move prices relative to current rents.⁵

While the literature to date has centered on these polar cases of perfect segmentation and zero segmentation, we introduce a framework that allows for intermediate levels of frictions, corresponding to an upward-sloping tenure supply curve as in Panel (c). In this case, a credit expansion shifting tenure 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 a single shock. However, one 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. To disentangle movements along the tenure supply curve from shifts in the curve, we need to discipline the slope of the tenure supply curve.

We do so empirically. To measure the slope of the tenure supply curve, we follow the simultaneous equations literature in using identified credit supply instruments that shift tenure demand but not tenure supply. With this slope in hand, the remainder of the paper constructs a structural model, calibrates it to match this estimated slope, and uses it to decompose the role of credit in the 2000s housing boom and bust.

Before moving on, it is worth noting that nothing in our argument is specific to credit or the 2000s cycle. The same logic we use here holds for *any* factor that influences the relative demand for homeownership in *any* time period. For instance, the slope of the tenure supply curve would affect the extent to which foreclosures impact prices due to reduced demand from foreclosed-upon households who are unable to obtain a mortgage (Guren and McQuade (2020)), the impact of a mortgage interest deduction on price-rent ratios (Steinberg and Rotberg (2024)), and the effect of property taxes on price-rent ratios.

2 Empirical Analysis

Motivated by the intuition in Section 1, our goal is to estimate the slope of the tenure supply curve, equal to the elasticity of the price-rent ratio to an identified tenure demand shock divided by the elasticity of the homeownership rate to that same shock. For the

⁵A shift in landlord discount rates for the same set of rental cash flows has a similar effect.

shock, we use changes in the supply of mortgage credit to households, which shifts housing demand as in Panel (c) of Figure 2. Estimating this slope by regressing the price-rent ratio on homeownership is problematic because these quantities can be affected by both supply and demand through a standard simultaneity argument, and because of measurement error. To address these issues, we seek instruments for credit supply, which isolate demand-side variation.⁶

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

2.1 Data

We construct a balanced annual panel spanning 1995 to 2017 at the core-based statistical area level, or at the metropolitan division level for those areas comprised of smaller divisions (we will henceforth refer to both types of areas as CBSAs). Our panel merges together data on house prices, rents, homeownership rates, credit, and other variables we use as 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 2.2 - 2.4. 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 Federal Housing Finance Administration (FHFA) indices in the Appendix.

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.⁷ It is available quarterly for 53 CBSAs beginning in 1989 and 62 CBSAs beginning in 1994.⁸

⁶It is possible that the instruments we use also affect our tenure supply curve. For instance, our main instrument shifts the supply of GSE mortgages in a local area. GSE mortgages are primarily given to owner-occupants, but about 5% of GSE mortgages are given to investors. If our instruments affect tenure supply as well, we will uncover a locus of equilibria rather than recovering the supply curve. We show in Section 6 that this is still useful for calibrating structural models.

⁷CBRE 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.

⁸Beginning our sample in 1995 allows us to use a fully balanced panel with the lagged price-rent ratio as a control. Appendix B presents alternative results using an unbalanced panel from 1990 to 2017.

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.

We use two main data sources for homeownership. First, we use data from the Census Housing and Vacancy Survey (HVS), which provides a panel of annual estimates of the homeownership rate at the CBSA level from 1986 to 2017. These data are noisy as they are obtained from a supplement to the Current Population Survey with only 72,000 households nationwide, and are also complicated by decennial changes in CBSA definitions. Rather than using an unbalanced panel with changing CBSA definitions, we drop CBSAs with significant changes in their definitions.⁹

Second, to address concerns about geographic coverage and noise in the HVS, we introduce a new measure of homeownership rates constructed directly from micro data. We briefly describe the construction of our new homeownership rate here with details available in Appendix C. We combine property ownership records from Zillow with individual address histories from Infutor. We standardize owner and occupant names and addresses, and assess whether a property is owner occupied based on whether any occupant shares a last name with any owner. We then aggregate across occupied units to create homeownership rate time series at the CSBA-level. To account for time trends in Infutor coverage, we adjust our time series to match low-frequency trends in homeownership in the Decennial Census and American Community Survey (ACS).

Appendix C.6 compares our data to the HVS, showing that our new data are much less volatile, match the ACS data more closely, and cover many more CBSAs and metropolitan divisions (390 in our data, compared to 75 in the HVS without significant CBSA changes). This greater coverage and reduced noise yield vastly improved statistical precision.

2.2 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 purchase by Fannie Mae and Freddie Mac. Since Fannie Mae

⁹Specifically, if the homeownership rate changes by more than 4 percentage points due to redefinition based on data from the two nearest decennial Censuses, we drop the CBSA. All missing CBSAs with TW data, which include several large CBSAs such as New York and Los Angeles, are included with our second homeownership data source.

and Freddie Mac have been found to offer subsidized interest rates (see e.g., Ambrose, LaCour-Little and Sanders (2004)), an increase in the CLL is thus a positive shock to the supply of mortgage credit for borrowers newly able to take advantage of it.

Although updates to the CLL are typically uniform across the country, LS argue 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. For example, an average of 8.3% 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 0.6% in El Paso. Our instrument exploits that a change in the CLL should have a bigger average effect in San Francisco than in El Paso.

We follow LS and define the instrument as $Z_{i,t} = ShareNearCLL_{i,t} \times \%ChangeInCLL_t$, where $ShareNearCLL_{i,t}$ is the fraction of HMDA mortgage originations in the prior year that are within 5 percent of the current year’s CLL. Since the CLL has occasionally varied by region, we use only changes in the national CLL to construct our instrument.¹⁰

We estimate the reduced form impulse response of an outcome variable of interest $Y_{i,t} \in \{PRR, HOR\}$, representing the price-rent ratio and homeownership rate, respectively, to a change in the instrument $Z_{i,t}$. We estimate these responses via the local projection approach of Jordà (2005), using ordinary least squares to estimate:

$$\log(Y_{i,t+k}) = \zeta_i + \psi_t + \beta_k^Y Z_{i,t} + \theta X_{i,t} + \epsilon_{i,t}, \quad (1)$$

for $k = 0, \dots, 4$ where β_k^Y is our impulse response of interest, ζ_i are location fixed effects, ψ_t are time fixed effects, and $X_{i,t}$ are controls. We include one lag of $Z_{i,t}$ and $\log(Y_{i,t})$ as controls to purge serial correlation. The CBSA fixed effects absorb any average differences across areas, including their supply elasticity, while the time fixed effects 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 our estimates are based purely on the share-shift variation in our instrument. Standard errors are clustered by CBSA, and the regressions are weighted by 2000 population.

The identifying assumption for this instrument 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 in the cross-section and that also varies with changes in the national CLL in the time series. Such a variable would contaminate our estimate of

¹⁰In 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. Using a locally-based CLL 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.

β with omitted variable bias. 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 Appendix B demonstrating that time-varying city characteristics such as employment and industry shares are not driving our results.

For our estimates using the HVS homeownership rate, we use a balanced panel of 41 CBSAs from 1995 to 2017 that have consistent HVS homeownership rates and TW rents for the full sample period for our analysis.¹¹ For our estimates using our new microdata-based homeownership rate, we use a balanced panel of all 62 CBSAs with TW rent data from 1995 to 2017. In Appendix B we explore an expanded sample of 390 CBSAs for which we have house prices but not rents and find consistent results.

Figure 3 plots the impulse responses of both PRR and HOR to the LS instrument, corresponding to β_k^{PRR} and β_k^{HOR} in equation (1). Panel (a) displays results for the HVS homeownership rate and corresponding sample. The price-rent ratio, shown in blue dots, has a hump-shaped response, peaking at 16.5 at the one-year horizon and then gradually declining, with strongly significant responses up to 2 years.¹² In contrast, the estimated increase in the homeownership rate, which is shown in red triangles, peaks at 2.7, and is statistically indistinguishable from zero at most horizons. Dividing the PRR and HOR point estimates in Panel (a) yields estimated tenure supply slopes of 6.0, 6.2, and 11.8 at the 0-2 year horizons.¹³


Because the homeownership rate coefficients are statistically indistinguishable from zero, the confidence interval for the slope is unbounded and we cannot statistically reject a slope of infinity (equivalent to perfect segmentation). To allow for easier visualization, we instead compute confidence intervals for the *inverse* ratio $\beta_k^{HOR} / \beta_k^{PRR}$, which we block bootstrap by CBSA. Panel (c) plots the resulting confidence intervals, which are tightly estimated for horizons of 0-2 years, before losing statistical power as the PRR coefficients decay and lose significance. Over the more precisely estimated 0-2 year horizons, we ob-

¹¹In practice, the CLL never adjusts downward, so it typically remains flat during housing downturns until prices pass their previous peak. Consequently, most of the identifying variation is obtained over the period 1992 to 2006 as the national CLL did not vary through the subsequent housing bust and rebound.

¹²Decomposing this result, this behavior of the price-rent ratio is due mostly to prices; rent growth in response to the Loutskina-Strahan instrument is small and statistically insignificant (Appendix Figure B.1). Our results for the effect of credit on house prices are qualitatively consistent with those found in the literature, such as Glaeser et al. (2012), Adelino et al. (2023), Favara and Imbs (2015), and Di Maggio and Kermani (2017), although it is difficult to compare our results quantitatively since we estimate reduced form responses to instruments.

¹³Since 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.

Figure 3: Loutskina-Strahan Instrument LP Impulse Responses



Notes: 95% confidence interval shown in bars. The figure shows panel local projection estimates of the response of the indicated outcomes to the LS instrument $ShareNearCLL_{i,t} \times \%ChangeInCLL_t$ as estimated using equation (1). Control variables include $ShareNearCLL_{i,t}$ and its lag and lags of the instrument and outcome variables, and regressions are weighted by 2000 population. Panels (a) and (b) show the price/rent and homeownership rate for the HVS and GG homeownership rates, respectively, with standard errors clustered by CBSA. Panels (c) and (d) show the inverse ratio $\beta_k^{PRR} / \beta_k^{HOR}$ for the HVS and GG homeownership rates, respectively, with standard errors block bootstrapped by CBSA.

serve upper bounds for the *inverse* ratio between 0.365 and 0.407, which correspond to *lower* bounds for the true ratio of between 2.5 and 2.7.

Panels (b) and (d) repeat this analysis using our new microdata-based homeownership rate. The effects on the price-rent ratio in Panel (b) differ from those in Panel (a) because they are estimated on a larger sample of CBSAs, but similarly peak at the 1-year horizon at 11.3, before remaining elevated through the 4-year horizon. In contrast, we observe very small estimated responses for the homeownership rate, equal to only 0.062 at the 1-year horizon and peaking at 0.290 at the 4-year horizon. Dividing these coefficients to obtain point estimates for the tenure supply slope yields values between 27.3 and 188.9 depending on the horizon. In addition, Panel (d) shows that the broader geographic sample and reduced noise in our new GG measure of the homeownership rate allows for

dramatically smaller standard errors. In particular, the upper bounds of our inverse ratio vary between 0.045 and 0.069 over the 0-2 year horizons, corresponding to lower bounds of the true ratio between 14.5 and 22.0. Extending this analysis out to the 3-year and 4-year horizon, we again observe some loss of precision, but obtain lower bounds for the true ratio of at least 4.5 at all horizons.

Overall, both sets of estimates display strong responses of the price-rent ratio that are substantially larger than the corresponding movements in the homeownership rate, indicative of strong frictions. In Section 4 we calibrate our model to match our estimates using the HVS homeownership rate, whose smaller estimated tenure supply slopes imply a more conservative impact of credit on house prices.

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

Our second approach follows DK in exploiting the OCC's 2004 preemption of state-level anti-predatory-lending laws (APLs) for national banks as an expansion of credit supply. APLs were implemented by states starting in 1999 to limit the terms of mortgages made to riskier borrowers, and their preemption allowed OCC-regulated banks to expand credit to riskier borrowers. DK exploit this quasi-experiment by comparing counties with different exposures to national banks that were regulated by the OCC before and after the change, an approach that we implement at the CBSA 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 mortgages 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 local projection approach as with our LS instrument, and instead follow DK in using a reduced-form event study:

$$\log(Y_{i,t}) = \zeta_i + \psi_t + \sum_{k \neq \tau} \beta_k Z_i \mathbf{1}_{t=k} + \theta X_{i,t} + \epsilon_{i,t}. \quad (2)$$

The coefficients β_k represent the reduced form effect of the instrument at each date 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 control for both of these variables directly in addition to including CBSA and year fixed effects ζ_i and ψ_t .¹⁴ We follow DK

¹⁴Our controls $X_{i,t}$ also include the lag of the outcome variable as well as all additional controls used by DK in their original analysis. The only control DK use that we do not use is a proprietary and confidential measure of the share of loans originated to subprime borrowers that is not crucial for their results.

Figure 4: Di Maggio-Kermani APL Preemption Reduced Form



Notes: The figure shows estimates of the cumulative sum from 2003 of β_k for each indicated year and outcome variable (price or GG-Microdata homeownership rate) estimated from equation (2), with the instrument being $Z_i = APL_{2004} \times OCC_{2003}$. 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 OCC-regulated lenders interacted with a dummy for post-2004, and the fraction of OCC-regulated lenders interacted with a dummy for APLs. All regressions are weighted by 2000 population and standard errors are clustered by CBSA, as in the original DK paper. 95% confidence intervals are shown in bars.

in estimating equation (2) using growth rates for $Y_{i,t}$, so that the outcome variable is either the log change in house prices or homeownership rates. We then obtain an impulse response in levels by cumulating the coefficients of interest from the base period to each indicated period and compute standard errors by the delta method.

The identifying assumptions are similar to a differences-in-differences approach: there must be parallel trends in the absence of treatment. DK provide support for this identifying assumption in their paper.

DK kindly provided us with their data set, and we directly use their data, 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 our new microdata-based homeownership rate. 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 2000 population, dropping states where the anti-predatory lending laws change in years other than the year of the OCC preemption (WI, IN, and MA), and clustering standard errors by CBSA. This yields a balanced panel of 370 CBSAs from 2001 to 2010.

Figure 4 displays our results for our new GG microdata-based homeownership rate.¹⁵

¹⁵Our results are similar to those found in Di Maggio and Kermani (2017), but differ for three reasons. First, we cumulate the β coefficients from the base period to the indicated period to obtain an IRF in levels

We observe a slight downward pre-trend prior to 2003 for house prices (which would work against our subsequent findings) and no pre-trend for homeownership. After 2004, house prices rise through 2007, while the response of homeownership rates is both statistically and economically insignificant. Taking the ratio of the coefficients for house price to those for homeownership, we obtain ratios between 20.6 and 27.1 from 2005 to 2007 and cannot reject ratios of infinity, consistent with our previous findings indicating strong frictions or segmentation. Block bootstrapping these ratios, as we did for our LS estimates, generates a lower bound for this ratio's 95% confidence interval of at least 2.2 in years 1-3. Appendix B shows results for the smaller and noisier HVS sample; unfortunately the confidence intervals are so wide that the results are uninformative.

2.4 Third Empirical Approach: Mian and Sufi (2019)


Our third approach follows MS in exploiting 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. MS build on evidence from Justiniano, Primiceri and Tambalotti (2022) of a sudden, sizable, and persistent expansion in PLS markets in late 2003 that persists until the crash. Following Nadauld and Sherlund (2013), MS argue that the PLS expansion had a larger effect on lending by mortgage lenders that rely on non-core deposits to finance mortgages, measured at the bank level as the ratio of non-core liabilities to total liabilities (NCL), and show that high NCL banks did in fact expand their lending more after following roughly parallel trends prior to 2002.

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 from equation (2) that we use for the DK instrument. We follow MS in weighting by 2000 housing units and including year and CBSA fixed effects, and cluster at the CBSA level. Our sample consists of a balanced panel of 258 CBSAs for which all of the data necessary is available.

Figure 5 displays our resulting β_k coefficients by year. We observe zero effect on house prices prior to 2002, followed then a hump-shaped impulse response peaking in 2006. By contrast, the homeownership rate rises by an economically and statistically insignificant

rather than in changes. Second, we estimate our results at the CBSA rather than county level. Third, we leave out one control for subprime share that they could not share with us due to data confidentiality. In Appendix Figure B.8, we replicate DK's results with the broader sample of cities and a non-cumulated IRF.

Figure 5: Mian-Sufi PLS Expansion Reduced Form



Notes: The figure shows estimates of the effect of a city's NCL share on the indicated outcome (price or homeownership rate) based on estimating equation (2) with the instrument being $Z_i = NCLShare_i^{2002}$ and 2002 being the base year. All standard errors are clustered by CBSA and all regressions are weighted by housing units as in the original MS paper. 95% confidence intervals are shown in bars.

amount. The ratios of the point estimates are large, between 15.2 and infinity depending on the horizon. Block bootstrapping these ratios yields a 95% confidence interval lower bound for this ratio of at least 2.1 in years 1-3, allowing us to again soundly reject a perfectly elastic tenure supply curve. As with MS, Appendix B shows that the DK instrument lacks sufficient power to estimate an informative slope using the HVS.

2.5 Summary

While all three of our empirical strategies have limitations, they all robustly find that house prices or price-rent ratios respond significantly more than homeownership rates, with implied ratios of the responses of at least 3.8 and lower bounds of the confidence interval of at least 2.1. We obtain these consistent results despite our instruments relying on distinct sources of variation and influencing different segments of the mortgage market, with the LS instrument affecting credit to prime borrowers and the DK and MS instruments affecting credit to more risky borrowers. With an empirical estimate of the slope in hand, we next introduce a model that we discipline with our empirical estimates, then use it to decompose the role of credit in the 2000s housing boom and bust.

3 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.

The model extends Greenwald (2018) by adding heterogeneity in the benefit to owning or renting for borrowers and landlords to allow for a fractional and time-varying home-ownership rate. The model’s key equilibrium conditions map directly into the supply and demand framework presented in Section 1. By adjusting the degree of heterogeneity in landlord ownership benefits, our framework nests models with a fixed homeowner-ship rate, with perfect arbitrage between renting and owning, and with any intermediate case, corresponding to Panels (a), (b), and (d) of Figure 2 respectively.

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 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). A fraction δ of housing must be replaced as maintenance each period.

In our baseline model, we assume that savers always demand a fixed quantity of housing \bar{H}_S or equivalently trade in a completely segmented housing market (e.g., live in different neighborhoods or occupy different quality tiers). This assumption reflects real world frictions due to indivisibility and variation in the quality and location of housing. We discuss these frictions and evaluate the robustness of our results to relaxing this strong assumption in Section 6. 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 $\hat{H}_t = \bar{H}_t - \bar{H}_S$.

To ensure that mortgage and housing flows are gradual, as in the data, a random fraction ρ of households are hit with moving shocks, which if they are owners leads them to sell and prepay their mortgages. Next, a random fraction η of moving households become “active,” making them eligible to buy housing and obtain new mortgages, which at equilibrium they will choose to do. Since movers include households that were previously

both owners and renters, calibrating η to equal the steady state borrower homeownership rate ensures that the per-buyer amount of purchased housing (i.e., house size) is similar to the per-seller amount of sold housing. This is important for a realistic application of our LTV and PTI constraints below.¹⁶ All borrower households, including non-movers, decide whether to purchase homes of equal size on an internal market. Our model's timing, discussed in detail in Appendix A, yields a symmetric allocation of owned housing across all borrower households with $\omega_{i,t}^B > \bar{\omega}_{B,t}$.

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

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

where c is nondurable consumption, and h is housing services. To nest typical specifications in the literature which model landlords as risk-neutral (see e.g., Kaplan et al. (2020)), we assume that landlords have linear utility, maximizing:

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

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 $(r^* + \nu)M^*$, $(1 - \nu)(r^* + \nu)M^*$, $(1 - \nu)^2(r^* + \nu)M^*$, et cetera. Mortgage loans are prepayable, with all active borrowers prepaying their mortgages each 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:

$$M_{i,t}^* \leq \theta^{LTV} p_t H_{i,t}^* \quad (3)$$

¹⁶For instance, if half of borrowers own housing, then half of moving borrowers are selling housing. If all movers then bought housing there would be twice as many buyers as sellers, so that each newly purchased house would be half the size of the sold ones, on average. These smaller houses would require smaller loans, artificially relaxing PTI limits. Setting $\eta = 1/2$ in this case reconciles purchased and sold house sizes.

$$M_{i,t}^* \leq \frac{(\theta^{PTI} - \kappa) y_{B,t} e_{i,t}}{r_{B,t}^* + \nu + \alpha}, \quad (4)$$

where p_t is the price of housing, $H_{i,t}^*$ is the new house size, $e_{i,t}$ is an idiosyncratic income shock generating variation in which constraint binds (see Greenwald (2018)), and κ 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 identical borrowers share a common valuation for housing, and identical landlords do the same, then whichever group has the higher valuation will own all the housing, leading to a homeownership rate of either 0% or 100%. To generate a fractional homeownership rate, we need to impose further heterogeneity in how agents value housing *within* at least one of these types. Our key modeling contribution in this paper is to introduce and carefully calibrate this within-type heterogeneity.

We impose this heterogeneity by assuming that agents receive an additional service flow (either positive or negative) from owning housing. Specifically, each borrower i receives surplus equal to $\omega_{i,t}$ times the market rent for their owned housing, where $\omega_{i,t}^B \sim \Gamma_{\omega,B}$ is drawn i.i.d. across borrowers and time. Symmetrically, if a landlord owns unit j of housing, he or she receives surplus equivalent to $\omega_j^L \sim \Gamma_{\omega,L}$ times the market rent for that unit. Because we perceive these benefits and costs as likely non-financial, particularly for borrowers, we rebate them lump-sum to households so they do not have any effect on the resource constraint in equilibrium.¹⁷ 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 ω_j^L are owned, and they are owned by the households with the largest $\omega_{i,t}^B$.¹⁸ We define the cutoff values as $\bar{\omega}_{B,t}$ and $\bar{\omega}_{L,t}$.

There are several forms of heterogeneity that map intuitively into this framework.

¹⁷Our current specification is isomorphic to one where instead of the ω^B shocks causing monetary gains and losses, borrowers instead have flow utility of the form:

$$u(c, h; \omega^B) = c^{1-\zeta} h^\zeta (1 + \omega^B \mathbf{1}_{own}),$$

where $\mathbf{1}_{own}$ is an indicator for homeownership and ω^B is drawn i.i.d. across households as before. This specification would also generate (5) as the optimality/indifference condition for the marginal homeowner. We choose to put ω^B as a monetary transfer in the budget because this transfer can be imposed in a parallel way for landlords, who do not have utility over housing services. Additionally, whether utility is increasing in ω^B above technically depends on whether $h > 1$, meaning that the sorting of owners by ω^B could hypothetically reverse depending on construction.

¹⁸Because of perfect risk sharing, the identities of the individual borrowers are irrelevant to the equilibrium. This means that our simplifying assumption that the ω^B shocks are i.i.d. rather than persistent is without loss of generality, as we could redraw the i indices i.i.d. each period without consequence.

On the borrower side, heterogeneity in the value of ownership stands in for household age, family composition, ability to make a down payment, and personal preference for ownership. On the landlord side, we believe the largest source of heterogeneity is on the suitability of different properties for rental (e.g., Halket, Nesheim and Oswald (2020)). For example, while urban multifamily units can be efficiently monitored and maintained as rental units, the depreciation and moral hazard concerns for renting a detached suburban or rural house may be more significant. As the homeownership rate rises, the marginal 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.

Borrower's Problem. The borrower maximizes expected lifetime utility subject to the borrowing constraints (3), (4), and their budget constraint:

$$\begin{aligned}
c_{B,t} \leq & \underbrace{(1-\tau)y_{B,t}}_{\text{after-tax income}} + \underbrace{\rho\eta \left(M_{B,t}^* - \pi^{-1}(1-\nu)M_{B,t-1} \right)}_{\text{net mortgage iss.}} - \underbrace{\pi^{-1}(1-\tau)X_{B,t-1}}_{\text{interest payment}} - \underbrace{\nu\pi^{-1}M_{B,t-1}}_{\text{principal payment}} \\
& - \underbrace{\rho p_t (\eta H_{B,t}^* - H_{B,t-1})}_{\text{net housing purchases}} - \underbrace{\delta p_t H_{B,t-1}}_{\text{maintenance}} - \underbrace{q_t (h_{B,t} - H_{B,t-1})}_{\text{rent}} \\
& + \underbrace{\left(\int_{\bar{\omega}_{B,t-1}} \omega d\Gamma_{\omega,B} \right) q_t \hat{H}_{t-1}}_{\text{owner surplus}} + \underbrace{T_{B,t}}_{\text{rebates}},
\end{aligned}$$

where $y_{B,t}$ is exogenous outside income, η is the share of active borrowers who choose to buy housing, q_t is the rental rate (i.e., the price of housing services). The threshold $\bar{\omega}_{B,t}$ is implicitly defined by the market clearing condition $1 - \Gamma_{\omega,B}(\bar{\omega}_{B,t}) = H_{B,t}/\hat{H}_t$, which ensures that the fraction of borrowers choosing to own (with $\omega_{i,t}^B > \bar{\omega}_{B,t}$) is equal to the fraction of borrower-inhabited housing owned by borrowers. Income is taxed at rate τ , while mortgage interest payments are tax deductible.

The borrower's state variables are the mortgage balance $M_{B,t}$, promised interest payment $X_{B,t}$, and owned housing $H_{B,t}$, where separate tracking of the interest payment $X_{B,t}$ is necessary due to the fixed-rate nature of the mortgages.¹⁹ The laws of motion are:

$$M_{B,t} = \underbrace{\rho\eta M_{B,t}^*}_{\text{new loans}} + \underbrace{(1-\rho)(1-\nu)\pi^{-1}M_{B,t-1}}_{\text{old loans}}$$

¹⁹The mortgage balance $M_{B,t}$ and interest payment $X_{B,t}$ must be separately tracked because under fixed-rate mortgages the current balance is not sufficient to determine the mortgage payment, which instead depends on the interest rate at the time the loan was originated. Introducing the additional state variable $X_{B,t}$ solves this problem.

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

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 6. The landlord maximizes expected lifetime utility subject to the budget constraint:

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

and the market clearing condition $\Gamma_{\omega,L}(\bar{\omega}_{L,t}) = H_{B,t}/\hat{H}_t$, which specifies that the fraction of properties not owned by landlords is equal to the share of borrower-inhabited housing owned by borrowers.

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

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

where the wedge Δ_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 $\Delta_t > 0$ 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. Following Davis and Heathcote (2007), Favilukis et al. (2017) and Kaplan et al. (2020), new housing is produced by competitive construction firms using nondurables Z_t and land L_t according to the technology:

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

\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 price of land permits. We assume for tractability that newly constructed housing has the same distribution of rental suitability (ω^L) as the existing stock.²⁰

Including construction in the model is 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.6). This is due to our focus on price-rent ratios rather than prices. Intuitively, new construction increases the quantity of both housing and housing services, which typically decreases both prices and rents in the same proportion, with minimal impact on the price-rent ratio.²¹

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 $(c_{S,t}, M_{B,t}^*)$, 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:²²

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

²⁰If builders could direct construction toward properties that are more or less suitable for rental, the tenure supply curve would be more elastic (flatter) in the long run than in the short run. For instance, a contraction in credit would induce a drop in prices in the short run, as households would be attempting to sell properties to landlords that are less and less suitable for renting. In the long run, directed construction could replace these properties with those more suitable for renting, allowing prices to recover and the homeownership rate to decrease. If this were the case, the response of house prices to credit shocks would weaken over time, while the response of the homeownership rate to credit shocks would strengthen over time. Because none of our empirically estimated homeownership responses in Section 2 strengthen significantly over time, we believe that this directed construction mechanism is not quantitatively strong.

²¹For instance, following an increase in housing demand, a model with perfectly inelastic construction supply would see a large increase in prices housing quantities (and rents) fixed, while a model with perfectly elastic construction supply would see a large increase in housing construction, decreasing rents. However, in both cases the price-rent ratio would rise, and in practice by similar amounts.

²²In 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.

3.1 Key Equilibrium Conditions

The key equilibrium conditions are the optimality conditions for borrower and landlord housing, which generate the tenure demand and supply curves:

$$p_t^{\text{Demand}}(H_{B,t}) = \frac{E_t \left\{ \Lambda_{B,t+1} \left[(1 + \bar{\omega}_{B,t})q_{t+1} + (1 - \delta - (1 - \rho)\mathcal{C}_{B,t+1})p_{t+1} \right] \right\}}{1 - \mathcal{C}_{B,t}} \quad (5)$$

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

where p_t^{Demand} is the price at which borrowers are willing to purchase quantity $H_{B,t}$, and p_t^{Supply} is the price at which landlords are willing to provide quantity $H_{B,t}$ to the market. By market clearing, we must have $H_{L,t} = \hat{H}_t - H_{B,t}$. The remaining equilibrium conditions are relegated to Appendix A.1.

Both schedules are based on standard asset pricing relations between an asset's price and 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 (6) 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 resale value of the housing net of depreciation. The demand schedule (5) 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).²³ 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 (5).

These equations map directly into the tenure demand and tenure supply schedules in the simple framework of Section 1. To see this, note that we can write $\bar{\omega}_{B,t} = \Gamma_{\omega,B}^{-1}(1 - HOR_{B,t})$ and $\bar{\omega}_{L,t} = \Gamma_{\omega,L}^{-1}(HOR_{B,t})$, where $HOR_{B,t} \equiv H_{B,t}/\hat{H}_t$ is the homeownership rate among borrowers (owner-occupied share of non-saver housing). As $HOR_{B,t}$ increases, $\bar{\omega}_{B,t}$ falls as the marginal household becomes less suited for ownership, generating a downward-sloping demand curve. At the same time, $\bar{\omega}_{L,t}$ rises with $HOR_{B,t}$ as the marginal property becomes more suitable for rental, generating an upward-sloping tenure supply curve. In equilibrium, $HOR_{B,t}$ adjusts so that $p_t^{\text{Demand}} = p_t^{\text{Supply}}$, and the market clears.

²³The exact definition is $\mathcal{C}_{B,t} \equiv \mu_{B,t} F_t^{LTV} \theta^{LTV}$ 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.

The dispersions of the $\Gamma_{\omega,B}$ and $\Gamma_{\omega,L}$ distributions map into the slopes of the tenure demand and tenure supply curves, respectively. The more dispersed are the ownership benefits, the more $\bar{\omega}$ moves for a given change in the homeownership rate (equivalently, a given change in $\Gamma_{\omega}(\bar{\omega})$). This leads to a larger change in price and a steeper slope. In contrast, a distribution with low dispersion will yield a more elastic curve, as agents share highly similar valuations. In the limit, zero and infinite dispersion correspond to perfectly elastic (horizontal) and perfectly inelastic (vertical) price schedules, respectively.

4 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 4.1, then calibrate $\sigma_{\omega,L}$ to directly match our empirical regressions in Section 4.2.

4.1 Main Calibration

Demographics and Preferences We use the 1998 Survey of Consumer Finances to obtain a borrower population share of $\chi_B = 0.626$ and an income share of $s_B = 0.525$.²⁴ For landlords, we consider the limit $\chi_L \rightarrow 0$ and assume that landlords do not receive labor income, instead subsisting entirely on their rental earnings.²⁵

A key preference parameter is the borrower’s discount factor, β_B , which determines the 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. Many borrowers choose to pay for PMI, while many do not, meaning that the typical borrower should be close to indifferent.²⁶ 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 1.75% of the loan, plus a spread of 80 basis points.²⁷

²⁴In 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. These groups would find it difficult to purchase a home without credit.

²⁵Because landlord utility is linear in consumption, assumptions about their income and consumption have essentially no impact on the results.

²⁶For example, 37.7% of Fannie Mae purchase loans required PMI over the 1999-2008 boom period (source Fannie Mae Single Family Data Set).

²⁷We 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

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	ζ	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)	$\mu_{\omega,L}$	0.861	Y	Avg. homeownership rate
Landlord het. (scale)	$\sigma_{\omega,L}$	9.336	Y	Empirical elasticities (Section 4)
Borr. het. (location)	$\mu_{\omega,B}$	0.211	Y	Borr. VTI (1998 SCF)
Borr. het. (scale)	$\sigma_{\omega,B}$	0.697	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>				
Share of active movers	η	0.496	Y	\bar{H}_B/\hat{H}
Moving rate	ρ	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)
PTI offset (other debt)	κ	0.08	N	Greenwald (2018)
PTI offset (taxes + insurance)	α	0.090%	N	Greenwald (2018)
Landlord LTV Limit	θ_L^{LTV}	0.000	N	No landlord credit

For the other preference parameters, we use a standard consumption weight parameter of $\zeta = 0.2$ on housing, following 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

80%. Goodman and Kaul (2017) show that the overall costs of the two forms of insurance are similar.

prices. This implies that while saver demand is fixed in the short run, it is at the correct “long run” equilibrium value. Last, we set the landlord discount rate β_L to be equal to β_B , which 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.

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 of these as logistic distributions 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 $\sigma_{\omega,L}$ controls the slope of the tenure supply curve and is calibrated to directly match our empirical estimates in Section 2. This is the most important parameter for our results, so its calibration procedure is described in detail separately in Section 4.2.

Next, we jointly set the average levels of borrower ownership utility ($\mu_{\omega,B}$) and landlord ownership utility ($\mu_{\omega,L}$) to match the the average ratio of home value to income among borrowers who own homes in the 1998 SCF, equal to 8.81 (quarterly), as well as 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.

Finally, we calibrate borrower ownership dispersion $\sigma_{\omega,B}$, which determines the slope of the demand curve, to match evidence from Berger, Turner and Zwick (2020), hereafter BTZ, on the impact of the First Time Home Buyer credit. BTZ estimate that this policy, which subsidized the purchase of housing by up to \$8,000, led 3.2% of eligible renters to switch to ownership from February 2009 to July 2010. Since \$8,000 is 3.68% of the average median sales price for US houses over the period 2009:Q1 to 2010:Q2 (source: US Census Bureau, FRED code MSPUS), we set $\sigma_{\omega,B} = 0.697$ so that exactly 3.2% of renters would switch from renting to owning if their housing purchases were given a 3.68% subsidy. It is worth noting that $\sigma_{\omega,B}$ is much less important for our results than $\sigma_{\omega,L}$ because expansions in credit shift the demand curve and then travel along the supply curve, making the supply slope far more influential than the demand slope (see Appendix Figure A.5 for sensitivity analysis for this parameter).

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. We set $\eta = 0.496$, so that the share of active borrowers is equal to the steady state share of borrowers that own, so that there are an equal number of moving home buyers and moving home sellers in steady state. For the remaining mortgage contract parameters, we follow Greenwald (2018), who provides a detailed calibration for this mortgage structure.

4.2 Calibration of Landlord Heterogeneity to Our Empirical Results

We calibrate the dispersion in the landlord’s ownership cost $\sigma_{\omega,L}$ — the key parameter governing the slope of the tenure supply curve and the response of house prices to a change in credit conditions — so that the model reproduces as closely as possible the estimated responses of the price-rent ratio and homeownership rate to our LS credit supply shock displayed in Panel (a) 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. We match results using the HVS homeownership rate as they imply flatter tenure supply slopes than the GG homeownership rate and thus represent a lower bound for the effect of credit on house prices.

Since a change in the CLL effectively changes the subsidy on new mortgages, the model analogue of the empirical regression is the log-linearized impulse response to a shock to the mortgage spread Δ_t , which we assume follows an AR(1).²⁸ We thus choose

²⁸In our setting, comparing regions that receive larger or smaller increases in the share of subsidized loans is isomorphic to computing the impulse response to a shock to the mortgage spread Δ_t for two reasons. First, we assume within-type risk sharing, which implies that a higher share of loans receiving a subsidy is equivalent to a change in the average interest rate on all loans. Second, lenders are local, so the GE response of local mortgage interest rates to a local credit shock is the same as the aggregate response of mortgage interest rates to an aggregate credit shock. Our assumption of local lenders is quantitatively unimportant; re-calibrating the model assuming fixed interest rates (i.e., set at the national level) delivers

$\sigma_{\omega,L}$ along with the size and persistence of the shock to the mortgage spreads $\Delta_{B,t}$ to minimize the squared error between the annualized model and data impulse responses scaled by the statistical uncertainty around the empirical estimates.²⁹ Formally, we minimize:

$$Q = \sum_{v \in \{PRR, HOR\}} \sum_{k=1}^4 \left(\frac{IRF_{v,k}^{Model} - \beta_{v,k}}{SE_{v,k}} \right)^2, \quad (7)$$

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 choose to match responses at horizons of 1 to 4 years. We exclude Year 0 because the shock in reality can arrive part way through the year, attenuating the coefficient, while the shock in the model has its full effect in Year 0, as if the shock always arrived at the very start of the year.


Our procedure estimates a landlord cost dispersion ($\sigma_{\omega,L}$) of 9.336, a mortgage spread shock persistence of 0.881, and a mortgage spread shock size of -2.626 bp quarterly, where a negative shock size captures that spreads fall due to the subsidy. 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 represent the perfectly inelastic and perfectly elastic tenure supply examples in Section 1, respectively. To isolate the role of our main parameter, we 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 dot-dashed 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 3.81 and 23.50 times that of the homeownership rate depending on the horizon. By comparison, the Full Segmentation model (blue dashed line) delivers a nearly identical path of the price-rent ratio, showing that our Benchmark estimates imply a high degree of segmentation. However, the Full Segmentation model by construction generates no increase in the homeownership rate. Last, the No Segmentation

extremely similar results.

²⁹The 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 $\sigma_{\omega,L}$.

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. We omit Year 0 estimates, which are not used in the calibration exercise. 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. Panels (a) and (b) correspond to Panel (a) of Figure 3, while Panels (c) and (d) correspond to Panel (c) of Figure 3. The shaded red area indicates the range of outcomes from the lower bound estimate of $\sigma_{\omega,L} = 1.594$, to the upper bound estimate of $\sigma_{\omega,L} = \infty$, equivalent to the Full Segmentation case.

model (yellow dashed line) 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 1.

We next compute a “credible set” for $\sigma_{\omega,L}$ displayed as the shaded area in Figure 6 Panel (c) that reflects our 95% confidence interval for the inverse tenure supply slope estimates in Figure 3 Panel (c).³⁰ The upper bound, targeting the tops of the confidence intervals, yields an estimate of $\sigma_{\omega,L} = 1.594$, while the lower bound, targeting the bottoms

³⁰To do so, we choose values of $\sigma_{\omega,L}$ to minimize the distance to the upper and lower ends of the 95% confidence interval:

$$Q_{UB} = \sum_{k=1}^4 \left(\frac{IRF_{IR,k}^{Model} - (\beta_{IR,k} + 1.96 \times SE_{IR,k})}{SE_{v,k}} \right)^2$$

of the confidence intervals, is most closely matched using the Full Segmentation case $\sigma_{\omega,L} \rightarrow \infty$. Panel (d) shows that our Benchmark calibration provides a close fit of the inverse ratio as well.³¹ Furthermore, the No Segmentation economy falls far outside of the credible set, with inverse ratios at least 1.7 times the upper bounds of our empirical confidence intervals over our estimation sample.

4.3 Additional Model Validation

The previous exercise directly calibrates our model to match the *relative* response of price-rent relative to homeownership following a shock. We next present two out-of-sample tests of the model’s ability to map changes in credit conditions into the correct *absolute* changes in house prices as an additional model validation.

First, we test the model’s implied response of house prices to a change in interest rates. To do so we compare our results in Figure 6 to the estimates of Adelino et al. (2023), who find that the semi-elasticity of house prices to contemporaneous annual mortgage interest rates falls between 1.2 and 9.1. Dividing the change in log house prices in Figure 6 by the change in mortgage rates at the same horizon yields a model semi-elasticity of 1.7 on impact, which then declines with the horizon. Since these values are influenced by our assumed construction elasticity, we also compute the semi-elasticity for the price-rent ratio, which is less influenced by construction. We find a semi-elasticity for the price-rent ratio of 2.2 on impact, which again declines with the horizon. This comparison shows that our model delivers house price responses to changes in interest rates within, but on the conservative side of, the range found to be empirically plausible by the literature.

Second, we validate the model’s implied house price response following a change in credit standards. For this we use the empirical estimates of Johnson (2020), measures the effect of differential PTI limit policy between Fannie Mae and Freddie Mac on local house

$$Q_{LB} = \sum_{k=1}^4 \left(\frac{IRF_{IR,k}^{Model} - (\beta_{IR,k} - 1.96 \times SE_{IR,k})}{SE_{v,k}} \right)^2,$$

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. We compute our credible set based on the 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.

³¹In principle, matching the separate price-rent ratio and homeownership rate IRFs in Figure 3 Panel (a) and matching the ratio responses in Panel (c) are both theoretically valid approaches. In practice, we find that the resulting calibrations are virtually identical, so this choice is inconsequential.

price growth. Johnson (2020) finds that the lagged share of Freddie Mac loans — a measure of the local exposure to this change in policy — predicts the share of borrowers with PTI ratios above 50% with coefficient -2.90, and predicts house price growth over the two quarters after the policy change with coefficient -1.87. Dividing these responses implies that a PTI limit policy that reduces the share of borrowers with PTI ratios above 50% by 1% reduces house prices by 0.645%. In Appendix A.4, we reproduce this experiment in our model by comparing our baseline paths of house prices and the high-PTI share in our boom experiment to those in a hypothetical area with a tighter PTI limit designed to mimic Freddie Mac’s policies. This exercise yields model-implied ratios between 0.557% and 0.717% depending on the set of shocks applied, all of which are extremely close to the empirical estimate.

In summary, these tests provide additional confidence that, despite its parsimonious structure and tractability, our model is capable of reproducing quantitatively realistic responses of house prices to changes in credit prices and standards.

5 Model Results

We now conduct a series of experiments to quantitatively assess the role that credit played in the 2000s housing boom using our calibrated model. Table 2 summarizes the results of these experiments, with robustness to alternative values of $\sigma_{\omega,L}$ in Appendix Table A.2.

Credit Expansion Experiments. To begin, we simulate a relaxation of credit standards and evaluate the model’s implications for the evolution of debt and house prices. Our baseline experiment, inspired by KMV and the broader literature, unexpectedly relaxes LTV limits from 85% to 99% and PTI limits from 36% to 65% in 1998 Q1, with both changes perceived as permanent. The new standards are left in place until 2007 Q1, at which time they unexpectedly and permanently revert to their original values. We compute our model responses as perfect foresight paths outside of these two unexpected changes.³²

The results of this experiment are shown in Figure 7. To highlight the role of land-lord heterogeneity, we also plot the responses for the polar No Segmentation (perfectly elastic supply, $\sigma_{\omega,L} = 0$) and Full Segmentation (perfectly inelastic supply, $\sigma_{\omega,L} \rightarrow \infty$) 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

³²To be precise, we trace out the fully nonlinear transitions back to the steady state of the model given the contemporaneous set of parameters in a setting where the households assume zero probability of the arrival of any aggregate shocks.

Table 2: Results, Boom Experiments


Experiment	Price-Rent	Homeown.	Loan-Inc.
Peak Data Increase	51.5%	3.3pp	72.2%
<i>Credit Relaxation (Share of Peak Data Increase)</i>			
Full Segmentation	36%	0%	53%
Benchmark	32%	17%	51%
Est. Lower Bound	22%	71%	45%
No Segmentation	-2%	180%	30%
<i>Credit Relaxation + Decline in Rates (Share of Peak Data Increase)</i>			
Full Segmentation	77%	0%	86%
Benchmark	70%	35%	82%
Est. Lower Bound	47%	136%	68%
No Segmentation	2%	306%	36%
<i>Removing Credit Relaxation from Full Boom (Share of Peak Data Increase)</i>			
Full Boom (Benchmark)	100%	100%	104%
No Credit Relaxation (Benchmark)	47%	73%	29%
Full Boom (No Segmentation)	100%	100%	111%
No Credit Relaxation (No Segmentation)	97%	-149%	53%
<i>Credit Relaxation + Decline in Rates: Extensions (Share of Peak Data Increase)</i>			
Landlord Credit	70%	27%	78%
Saver Demand	51%	17%	102%

Notes: This table summarizes the results from the various nonlinear transition experiments in Sections 5 and 6. “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 residential mortgages (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.


for 32% 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 22%.³³ This stands in sharp contrast to the No Segmentation model, where the same credit relaxation explains -2% of the peak growth in price-rent ratios, as landlords completely absorb the increase in demand,

³³To be consistent with the data, the price-rent ratio is calculated as the ratio of the total value of housing to total housing services: $p_t \bar{H}_t / (q_t \hat{H}_t + q_{S,t} \bar{H}_S)$. The denominator includes imputed rents for all homeowners including savers, for whom the total imputed rent is the quantity of saver housing \bar{H}_S times the saver marginal rate of substitution $q_{S,t} = u_{S,t}^h / u_{S,t}^c$. Computing the price-rent ratio using only borrower housing and rents only leads to a more gradually rising price-rent ratio but a very similar peak response.

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 4, 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 $\sigma_{\omega,L} = 1.594$, 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.

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 36% of the observed rise in price-rent ratios.

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% vs. 30% 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 2pp 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 an additional increase in house prices. Combining the relaxation in LTV and PTI limits and the fall in rates can explain 70% of the observed rise in price-rent ratios and 82% of the rise in loan-to-income ratios, shown in Figure 7 Panel (b). These results again stand in contrast to the 2% and 36% 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.

At the estimated lower bound level of segmentation, we find that the combination of low rates and loose credit standards would explain 47% of the observed rise in the price-rent ratio, and 136% of the observed rise in homeownership. Thus, after accounting for uncertainty in our key moments, we can reject that credit played a small role in the rise of house prices during the boom, but cannot reject that credit was the primary factor driving the rise in homeownership.


Macroprudential Policy Experiments. We next study the potential effect of macroprudential policy on the boom-bust cycle in our various economies. To do so, we introduce a set of shocks that can fully explain the boom-bust for price-to-rent ratios and homeownership rates. We then compare these responses to counterfactual experiments where the same set of shocks arrive, but credit standards do not loosen due to a policy intervention.

To begin, we need to supplement our credit expansion experiment in Panel (b) of Figure 7 with additional shocks to fully explain the rise in price-rent ratios and homeownership rates over the boom period. These represent non-credit factors such as overoptimistic house price expectations that are widely believed to have played a major role in driving the boom (see e.g., Kaplan et al. (2020) and Chodorow-Reich, Guren and McQuade (2023)). To implement these, we follow the intuition in Section 1 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 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.³⁴

Last, we incorporate two features that help to generate a more realistic housing bust:

³⁴This 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.

Figure 8: Full Boom Experiment



Notes: Plots display perfect foresight paths. The “Benchmark” calibrates $\sigma_{\omega,L}$ to match our empirical IRFs as in Section 4, 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.

a further 3pp 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.³⁵ These additional assumptions are only relevant following the bust and have no impact on the model’s boom dynamics.

The resulting transition paths are plotted in Panel (a) of Figure 8. Overall, the model generates a good fit of the dynamics of the boom and bust, with two main exceptions: (i)

³⁵This is best interpreted as tight credit score standards preventing a fraction of the population from obtaining credit at all, rather than a decline in maximum LTV and PTI ratios at the intensive margin.

house prices jump in the model rather than adjust sluggishly in the data, which is typical for models lacking frictions to generate price momentum; 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 potential impact of macroprudential policy, we compute a counterfactual economy that features the same shocks to mortgage rates, tenure demand, and tenure supply (Δ_t , $\mu_{\omega,B}$, and $\mu_{\omega,L}$), but does not allow a relaxation of credit standards (θ^{LTV} , θ^{PTI}). The resulting responses (labeled “Bench. + Macropru.”) show that this policy would have been highly effective at dampening the housing cycle, reducing the overall rise in price-rent ratios by 53% of their actual rise and in loan-to-income ratios by 71% of their actual rise. These shares explained are larger than the those computed when relaxing credit standards in isolation (32% and 51%, respectively) because loose credit amplifies the role of non-credit demand factors like expectations.³⁶

Panel (b) repeats this quantitative exercise in the frictionless No Segmentation model, once again adding supply and demand shocks to match the boom in price-rent ratios and homeownership, then computing a “No Seg. + Macropru.” counterfactual that removes the relaxation of credit standards from the resulting set of shocks. Macroprudential policy is vastly less effective in the absence of segmentation, reducing the increase in price-rent ratios by only 3%. 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 47%. Thus, the ability of macroprudential policy to effectively stem excessive growth in credit depends crucially on the slope of the tenure supply curve, demonstrating the importance of our new moment for calibration.

Summary. Our calibrated model attributes an important role for credit conditions in explaining the housing and credit cycles observed in the 2000s boom-bust. We find that the fraction of the observed rise in house prices explained by the relaxation of credit standards is between 32% (when credit standards are relaxed in isolation) and 53% (when credit standards are relaxed alongside other contemporaneous shocks). These responses are much closer to those obtained under full segmentation than under a frictionless model with no segmentation, indicating the presence of strong frictions.

³⁶The intuition behind this finding, discussed at length in Greenwald (2018), is that PTI limits begin to bind as prices rise relative to incomes. Since it is difficult or costly for most households to produce large amounts of cash for a housing purchase, these binding limits constrain housing demand, even when borrowers are very optimistic. Thus, the impact of expectations on house prices is much larger when credit is simultaneously relaxed, compared to when it remains tight.

6 Model Extensions

In this section we test the robustness of our results by relaxing two important model assumptions: that landlords do not use credit, and that savers have fixed housing demand.

Landlord Credit. To parallel the existing literature, our baseline model assumes that landlords do not use credit. While this assumption makes the economics of the model transparent, it is clearly an abstraction, as many landlords both use credit and face financial constraints in reality. To address this, Appendix A.2 extends the model to allow landlords to purchase rental properties using credit. We assume that landlords borrow from savers using an identical technology to that used by borrower households, with landlord borrowing is limited by an LTV limit of 65% (a standard constraint for multi-family construction loans) and no PTI limit.

Under these assumptions, a shock that influences credit conditions for all borrowers (including landlords) now generates an upward shift in both the demand and supply curves, similar to the case displayed in Panel (d) of Figure 2. Holding the calibration fixed, this implies a larger response of the price-rent ratio and a smaller (or even negative) response of the homeownership rate in response to credit shocks. Recalibrating the model to match our empirical results, however, requires a somewhat flatter slope of the tenure supply curve. The recalibrated model produces largely similar results when paired with the shift in tenure supply — an outcome that can be visualized intuitively by comparing Panels (c) and (d) of Figure 2.

Appendix Figure A.1 reproduces the “Credit Relaxation + Decline in Rates” experiment from Figure 7 Panel (b) in the extended model with landlord credit under the assumption that landlords receive the same change in the mortgage rate and experience a similar LTV relaxation as households. Our results, summarized in the bottom panel of Table 2, display a nearly identical response of the price-rent ratio compared to the Benchmark model (each explaining 70% of the observed increase), and a smaller rise in the homeownership rate (explaining 27% vs. 35% of the observed increase). This occurs as an expansion of landlord credit leads landlords to compete more aggressively with households for properties, increasing prices and reducing the net quantity sold to households during the boom. The close similarity between these results and those of our baseline model provides reassurance that our baseline model’s simplifying assumption that credit shocks do not shift tenure supply does not lead to large bias in our results.

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 much like landlords do in the baseline model.

To evaluate the extent to which saver housing demand attenuates the role of credit quantitatively, Appendix A.3 extends the model to allow savers to freely trade housing. Because savers’ demand schedule is disconnected from credit conditions, allowing savers to act as buyers of housing effectively flattens the price schedule that borrowers face, leading to smaller responses of house prices and price-rent ratios to credit shocks. This is similar to a flattening of the tenure supply curve, with the exception that since savers in the model are also homeowners, transactions between borrower and saver households do not affect the homeownership rate.

Appendix Figure A.3 displays the results of the “Credit Relaxation + Declining Rates” experiment from Figure 7 Panel (b) in our saver demand extension, with the results again summarized in the bottom panel of Table 2. The results show that allowing for flexible saver demand does dampen the impact of credit on house prices. However, even in the extended model, credit changes are still able to explain 51% of the observed rise in the price-rent ratio, implying that this mechanism does not overturn our main results on the importance of credit in driving the 2000s boom-bust cycle.³⁷

We consider this saver extension to be an extreme lower bound on the strength of credit on house prices, for two reasons. First, we impose no heterogeneity (ω shocks) on the savers, meaning that the slope of saver demand is driven solely by diminishing marginal utility from housing, which minimizes the slope of the saver price schedule. Second, 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, further limiting trade across types. For instance, many properties purchased by constrained borrowers during the boom would not be appealing to wealthy savers as second homes during the bust, and certainly could not be absorbed into their primary homes as would be implied by the model. Indeed, Landvoigt et al. (2015) show that while changes in demand can ripple

³⁷This result is somewhat parameter dependent, as a specification in which savers had linear utility over housing would feature an effectively flat price schedule, neutralizing any impact on the price-rent ratio. However, we consider such a parameterization as unrealistic compared to our specification of log preferences, which are much more consistent with the stability of expenditure shares on housing documented by Davis and Ortalo-Magné (2011) and Piazzesi, Schneider and Tuzel (2007).

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.

7 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 instruments, 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 32% and 53% 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 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, Manuel, Antoinette Schoar, and Felipe Severino**, “Credit Supply and House Prices: Evidence From Mortgage Market Segmentation,” 2023. NBER Working Paper 31769.
- Ambrose, Brent W, Michael LaCour-Little, and Anthony B Sanders**, “The Effect of Conforming Loan Status on Mortgage Yield Spreads: A Loan Level Analysis,” *Real Estate Economics*, 2004, 32 (4), 541–569.
- Andersen, Steffen, John Y Campbell, Kasper Meisner Nielsen, and Tarun Ramadorai**, “Inattention and Inertia in Household Finance: Evidence From the Danish Mortgage Market,” 2014. Unpublished Paper, Available at SSRN, No. 2463575.
- Berger, David, Nicholas Turner, and Eric Zwick**, “Stimulating Housing Markets,” *The Journal of Finance*, 2020, 75 (1), 277–321.
- Campbell, Jeffrey R and Zvi Hercowitz**, “The Role of Collateralized Household Debt in Macroeconomic Stabilization,” 2005. National Bureau of Economic Research Working Paper No. 11330.
- Chodorow-Reich, Gabriel, Adam M. Guren, and Timothy J. McQuade**, “The 2000s Housing Cycle With 2020 Hindsight: A Neo-Kindlebergerian View,” *The Review of Economic Studies*, 2023, *Forthcoming*.
- Davis, Morris A and François Ortalo-Magné**, “Household Expenditures, Wages, Rents,” *Review of Economic Dynamics*, 2011, 14 (2), 248–261.
- **and Jonathan Heathcote**, “The Price and Quantity of Residential Land in the United States,” *Journal of Monetary Economics*, 2007, 54 (8), 2595–2620.
- Eggertsson, Gauti B and Paul Krugman**, “Debt, Deleveraging, and the Liquidity Trap: A Fisher-Minsky-Koo Approach,” *The Quarterly Journal of Economics*, 2012, 127 (3), 1469–1513.
- Favara, Giovanni and Jean Imbs**, “Credit Supply and the Price of Housing,” *American Economic Review*, 2015, 105 (3), 958–992.
- Favilukis, Jack, Sydney C Ludvigson, and Stijn Van Nieuwerburgh**, “The Macroeconomic Effects of Housing Wealth, Housing Finance, and Limited Risk Sharing in General Equilibrium,” *Journal of Political Economy*, 2017, 125 (1), 140–223.
- Foote, Christopher L, Kristopher S Gerardi, and Paul S Willen**, “Why Did So Many People Make So Many Ex Post Bad Decisions? The Causes of the Foreclosure Crisis,” in Alan S. Blinder, Andrew W. Lo, and Robert M. Solow, eds., *Rethinking the Financial Crisis*, New York: Russell Sage Foundation, 2012, pp. 136–186.
- Garriga, Carlos and Aaron Hedlund**, “Mortgage Debt, Consumption, and Illiquid Housing Markets in the Great Recession,” *American Economic Review*, 2020, 110 (6), 1603–1634.
- **and —**, “Housing Finance, Boom-Bust Episodes, and Macroeconomic Fragility,” 2022. Working Paper, Federal Reserve Bank of St. Louis and University of Missouri.
- **, Finn E Kydland, and Roman Šustek**, “Mortgages and monetary policy,” *The Review of Financial Studies*, 2017, 30 (10), 3337–3375.

- , **Rodolfo Manuelli, and Adrian Peralta-Alva**, “A Macroeconomic Model of Price Swings in the Housing Market,” *American Economic Review*, 2019, 109 (6), 2036–2072.
- Gete, Pedro and Michael Reher**, “Mortgage Supply and Housing Rents,” *Review of Financial Studies*, 2018, 31 (12), 4884–4911.
- Ghent, Andra**, “Infrequent Housing Adjustment, Limited Participation, and Monetary Policy,” *Journal of Money, Credit and Banking*, 2012, 44 (5), 931–955.
- Glaeser, Edward L, Joshua D Gottlieb, and Joseph Gyourko**, “Can Cheap Credit Explain the Housing Boom?,” in “Housing and the Financial Crisis,” University of Chicago Press, 2012, pp. 301–359.
- Goodman, Laurie and Karan Kaul**, “Sixty Years of Private Mortgage Insurance in the United States,” 2017. Research REport, Urban Institute.
- Greenwald, Daniel**, “The Mortgage Credit Channel of Macroeconomic Transmission,” 2018. MIT Sloan Research Paper No. 5184-16.
- Guren, Adam and Timothy McQuade**, “How Do Foreclosures Exacerbate Housing Downturns?,” *Review of Economic Studies*, 2020, 87 (3), 1331–1364.
- Guren, Adam M, Arvind Krishnamurthy, and Timothy J McQuade**, “Mortgage design in an equilibrium model of the housing market,” *The Journal of Finance*, 2021, 76 (1), 113–168.
- Halket, Jonathan, Lars Nesheim, and Florian Oswald**, “The housing stock, housing prices, and user costs: The roles of location, structure, and unobserved quality,” *International Economic Review*, 2020, 61 (4), 1777–1814.
- Huo, Zhen and Jose-Victor Rios-Rull**, “Financial Frictions, Asset Prices, and the Great Recession,” 2016. Working Paper, Federal Reserve Bank of Minneapolis.
- Iacoviello, Matteo**, “House Prices, Borrowing Constraints, and Monetary Policy in the Business Cycle,” *American Economic Review*, 2005, pp. 739–764.
- and **Stefano Neri**, “Housing Market Spillovers: Evidence From an Estimated DSGE Model,” *American Economic Journal: Macroeconomics*, 2010, 2 (2), 125–164.
- Johnson, Stephanie**, “Mortgage Leverage and House Prices,” 2020. Working Paper, Rice University.
- Joint Center for Housing Studies of Harvard University**, “The State of the Nation’s Housing 2020,” Technical Report 2020.
- Jordà, Òscar**, “Estimation and Inference of Impulse Responses by Local Projections,” *American Economic Review*, 2005, pp. 161–182.
- Justiniano, Alejandro, Giorgio E Primiceri, and Andrea Tambalotti**, “Household Leveraging and Deleveraging,” *Review of Economic Dynamics*, 2015, 18 (1), 3–20.
- , —, and —, “Credit Supply and the Housing Boom,” *Journal of Political Economy*, 2019, 127 (3), 1317–1350.

- , **Giorgio E. Primiceri**, and **Andrea Tambalotti**, “The Mortgage Rate Conundrum,” *Journal of Political Economy*, 2022, 130 (1), 121–156.
- Kaplan, Greg, Kurt Mitman, and Giovanni L. Violante**, “The Housing Boom and Bust: Model Meets Evidence,” *Journal of Political Economy*, 2020, 128 (9), 3285–3345.
- Kiyotaki, Nobuhiro, Alexander Michaelides, and Kalin Nikolov**, “Winners and Losers in Housing Markets,” *Journal of Money, Credit and Banking*, 2011, 43 (2-3), 255–296.
- Landvoigt, Tim, Monika Piazzesi, and Martin Schneider**, “The Housing Market(s) of San Diego,” *American Economic Review*, 2015, 105 (4), 1371–1407.
- Liu, Zheng, Pengfei Wang, and Tao Zha**, “Land-Price Dynamics and Macroeconomic Fluctuations,” *Econometrica*, 2013, 81 (3), 1147–1184.
- , —, and —, “A Theory of Housing Demand Shocks,” 2019. National Bureau of Economic Research Working Paper No. 25667.
- Loutskina, Elena and Philip E. Strahan**, “Financial Integration, Housing, and Economic Volatility,” *Journal of Financial Economics*, 2015, 115 (1), 25–41.
- Maggio, Marco Di and Amir Kermani**, “Credit-Induced Boom and Bust,” *The Review of Financial Studies*, 2017, 30 (11), 3711–3758.
- Mian, Atif and Amir Sufi**, “Credit Supply and Housing Speculation,” *Review of Financial Studies*, 2022, 35 (2), 680–719.
- Monacelli, Tommaso**, “Optimal Monetary Policy With Collateralized Household Debt and Borrowing Constraints,” in “Asset Prices and Monetary Policy,” University of Chicago Press, 2008, pp. 103–146.
- Nadauld, Taylor D and Shane M. Sherlund**, “The Impact of Securitization on the Expansion of Subprime Credit,” *Journal of Financial Economics*, 2013, 107 (2), 454–476.
- Piazzesi, Monika, Martin Schneider, and Selale Tuzel**, “Housing, Consumption and Asset Pricing,” *Journal of Financial Economics*, 2007, 83 (3), 531–569.
- Rognlie, Matthew, Andrei Shleifer, and Alp Simsek**, “Investment Hangover and the Great Recession,” *American Economic Journal: Macroeconomics*, 2018, 10 (2), 113–53.
- Saiz, Albert**, “The Geographic Determinants of Housing Supply,” *The Quarterly Journal of Economics*, 2010, 125 (3), 1253–1296.
- Steinberg, Joseph B. and Shahar Rotberg**, “Mortgage Interest Deductions? Not a Bad Idea After All,” *Journal of Monetary Economics*, 2024, *Forthcoming*.

Online Appendix

A Model Appendix

A.1 Model Details

This section presents details on model timing, the full set of equilibrium conditions, and the definition of equilibrium.

Model Timing. The timing of the borrower's problem proceeds as follows:

1. A random fraction ρ of borrowers receive moving shocks, which leads them to sell their housing in the main (inter-type) market at price p_t and prepay their mortgages.
2. A fraction ν of moving borrowers become active, allowing them to buy new housing in the main (inter-type) housing market at price p_t and obtain new mortgages.
3. All borrower-occupied housing (i.e., housing owned by either the borrower or the landlord) is divided into equal-sized units.
4. Borrowers observe their value of $\omega_{i,t}^B$ and choose whether or not to own by buying housing on the internal borrower market at price \tilde{p}_t per unit of housing.
5. Borrowers choose their nondurable and housing services consumption.

We have chosen this particular timing structure to overcome four challenges. First, if all households were allowed to participate in the inter-type mortgage and housing markets each period, allocations of housing and mortgages would adjust very quickly, which is both at odds with reality (see e.g., Andersen, Campbell, Nielsen and Ramadorai (2014)) and can make the model responses unstable. To address this, we assume that only a fraction ρ of households receive moving shocks allowing them to update their housing and mortgage allocations.

Second, since only a fraction of moving households were previously owners, if we allowed all moving households to buy housing, the typical size of the purchased property would be much smaller than the typical size of the properties being sold. For instance, if half of borrowers own housing (approximately true in our steady state), then half of moving borrowers are selling housing. If all movers then bought housing there would be twice as many buyers as sellers, implying that each newly purchased house would be half the size of the sold ones, on average. Purchasing smaller properties would artificially

relax the PTI limit, which binds with the value of the property is large compared to the borrower's income. Imposing that only a fraction η of mover households are able to buy, calibrated so that η is equal to the steady state fraction of borrowers who own, ensures that the sizes of purchased and sold properties are similar, avoiding this problem.

Third, without an internal market, housing would not be allocated to the highest ω^B borrowers at any given time. As a result, instead of our tenure demand and tenure supply schedules relating the price-rent ratio to the overall homeownership rate, we would obtain schedules relating the price-rent ratio to the homeownership rate among moving households only. While this is a reasonable specification, and we have computed results under this assumption, it is further from our intuitive motivation, and delivered much more sluggish responses of homeownership that were an inferior fit for the data. Our imposition of an internal borrower market addresses this issue by allowing housing to be reallocated to all borrowers with the highest ω^B values, restoring the relation between house prices and the overall homeownership rate.

Fourth, without further frictions, borrowers with different values of ω^B would buy different amounts of housing, making the problem much more complex. To keep the model parsimonious, we assume that housing is separated into equal-size blocks on the internal market, yielding a simple indifference condition that pins down house prices.

Borrower's Problem. First, we consider the problem of the fraction $\rho\eta$ of borrowers who are eligible to purchase housing on the inter-type market decide how much housing to purchase, to be sold later on the internal borrower market. The optimality condition for purchasing housing on this market at price p_t is

$$p_t = \frac{\tilde{p}_t}{1 - C_{B,t}} \quad (\text{A.1})$$

where \tilde{p}_t is the price of housing on the internal market. Intuitively, the numerator (the internal market price) reflects the value of a property that must be purchased outright and cannot be used as mortgage collateral. The denominator accounts for the marginal value of using the property as collateral to obtain mortgage credit, increasing p_t . The collateral value term $C_{B,t}$ is in turn defined by

$$C_{B,t} = \mu_{B,t} F_{B,t}^{LTV} \theta_{B,t}^{LTV}.$$

where $\mu_{B,t}$ is the multiplier on the borrowing constraint, $F_{B,t}^{LTV}$ is the fraction of borrowers who are LTV-constrained (see Greenwald (2018) for a full derivation of this expression).

The LTV-constrained share is in turn defined by:

$$F_{B,t}^{LTV} = \Gamma_e(\bar{e}_t)$$

$$\bar{e}_t = \frac{\theta^{LTV} p_t H_{B,t}^* (r_{B,t}^* + \nu + \alpha)}{(\theta^{PTI} - \omega) y_{B,t}}.$$

In the second stage, all borrower-inhabited housing is divided into equal portions, including borrower-owned housing. Borrowers then draw their owner surplus shocks $\omega_{i,t}^B$. For the market to clear, the fraction of borrowers who choose to own by buying on the internal market ($1 - \Gamma_{\omega,B}(\bar{\omega}_{B,t})$) must equal the share of borrower-inhabited housing that is owner-occupied ($H_{B,t}/\hat{H}_t$). The price of housing on the internal market then adjusts so that the marginal borrower is indifferent at equilibrium:

$$\tilde{p}_t = E_t \left\{ \Lambda_{B,t+1} \left[\bar{\omega}_{B,t} + q_{t+1} - \delta p_{t+1} + (1 - \rho) \tilde{p}_{t+1} + \rho p_{t+1} \right] \right\}. \quad (\text{A.2})$$

Equation (A.2) specifies that for the marginal borrower who buys housing, 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), net of the maintenance expense, plus the continuation value. With probability $1 - \rho$ the borrower will only be able to sell the property in next period's internal market at price \tilde{p}_t , but with probability ρ the borrower will receive a moving shock and be able to sell housing at the inter-type market price p_t , which is higher as it allows housing to either be sold to landlords or be used to collateralize new mortgages. Substituting in the relation $\tilde{p}_t = (1 - \mathcal{C}_{B,t})p_t$, which follows directly from (A.1) and manipulating this expression yields (5).

The borrower's optimality conditions for housing services ($h_{B,t}$) is

$$(h_{B,t}) : \quad q_t = (u_{B,t}^h / u_{B,t}^c), \quad (\text{A.3})$$

which sets the rent equal to the marginal rate of substitution between housing services and consumption.

The optimality condition for new mortgage debt for each active borrower who buys ($M_{B,t}^*$) is:

$$(M_{B,t}^*) : \quad 1 = \Omega_{M,t}^B + r_{j,t}^* \Omega_{X,t}^B + \mu_{B,t}, \quad (\text{A.4})$$

where $\bar{r}_{B,t-1} = X_{B,t-1}/M_{B,t-1}$ is the average rate on existing debt, and where the marginal

continuation cost of principal balance $\Omega_{M,t}^B$ and of interest payments $\Omega_{X,t}^B$ satisfy:

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

Equation (A.4) 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).

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

$$\begin{aligned}(B_t) : \quad & 1 = R_t E_t \left[\pi^{-1} \Lambda_{S,t+1} \right] \\ (M_{B,t}^*) : \quad & 1 = Q_{M,t}^S + r_{B,t}^* Q_{X,t}^S\end{aligned}$$

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

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

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

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

A.2 Extension: Landlord Credit

When landlords use credit, we impose that their problem becomes symmetric to the borrower's, with the same random selection to move (refinance their mortgages) and become active buyers. In this case we set the share of active buyers to $1 - \eta$, which is equal to the steady state share of borrower-inhabited properties owned by landlords. However, because landlords do not face PTI limits, we note that this assumption is not restrictive, and alternative assumptions about exactly which landlords are able to buy housing and in what quantities would lead to similar results, as LTV limits are linear in the amount of

housing purchased. When using credit, the landlord's budget constraint becomes:

$$\begin{aligned}
c_{L,t} \leq & \underbrace{(1-\tau)y_{L,t}}_{\text{after-tax income}} + \underbrace{\rho\eta \left(M_{L,t}^* - \pi^{-1}(1-\nu)M_{L,t-1} \right)}_{\text{net mortgage iss.}} - \underbrace{\pi^{-1}(1-\tau)X_{L,t-1}}_{\text{interest payment}} - \underbrace{\nu\pi^{-1}M_{L,t-1}}_{\text{principal payment}} \\
& - \underbrace{\rho\eta p_t (\eta H_{L,t}^* - H_{L,t-1})}_{\text{net housing purchases}} - \underbrace{\delta p_t H_{L,t-1}}_{\text{maintenance}} - \underbrace{q_t (h_{L,t} - H_{L,t-1})}_{\text{rent}} \\
& + \underbrace{\left(\int_{\bar{\omega}_{L,t-1}} \omega d\Gamma_{\omega,L} \right) q_t \hat{H}_{t-1}}_{\text{owner surplus}} + \underbrace{T_{L,t}}_{\text{rebates}},
\end{aligned}$$

while the landlord's laws of motion are:

$$\begin{aligned}
M_{L,t} &= \rho(1-\eta)M_{L,t}^* + (1-\rho)(1-\nu)\pi^{-1}M_{L,t-1} \\
X_{L,t} &= \rho(1-\eta)r_{L,t}^*M_{L,t}^* + (1-\rho)(1-\nu)\pi^{-1}X_{L,t-1} \\
H_{L,t} &= \rho(1-\eta)H_{L,t}^* + (1-\rho)H_{L,t-1}.
\end{aligned}$$

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 indifference condition for the marginal property now becomes

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

where $\mathcal{C}_{L,t} = \mu_{L,t} F_{L,t}^{LTV} \theta_{L,t}^{LTV}$ is defined analogously to the borrower case. The optimality condition for new credit issuance ($M_{L,t}^*$) becomes

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

The fixed point conditions that pin down the marginal continuation costs of debt are defined by:

$$\begin{aligned}
\Omega_{M,t}^L &= E_t \left\{ \Lambda_{L,t+1} \pi^{-1} \left[\nu + (1-\nu) \left(\rho_{L,t+1} + (1-\rho_{L,t+1}) \Omega_{M,t+1}^L \right) \right] \right\} \\
\Omega_{X,t}^L &= E_t \left\{ \Lambda_{L,t+1} \pi^{-1} \left[(1-\tau) + (1-\nu)(1-\rho_{L,t+1}) \Omega_{X,t+1}^L \right] \right\},
\end{aligned}$$

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 terms 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, so $\theta_L^{LTV} = 0.65$ and $\theta_L^{PTI} = \infty$. This implies $F_{L,t}^{LTV} = 1$.

Calibrating the model to match our empirical IRFs as in Section 4 requires mapping the identified LS shock into the model not only for borrowers but also for landlords. This is more subtle than for borrowers, for whom all loans are affected by the shock, because only single family rental properties and multifamily rental properties with fewer than 5 units are eligible for GSE financing and thus affected by changes in the conforming loan limit. 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. Beyond this, the calibration is the same as in the main text.

Results. We first solve this extension using our Benchmark value of $\sigma_{\omega,L}$. Figure A.1 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. Summary statistics are 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, implying that credit standards for landlords and households are relaxed to a similar degree.

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 75% of the rise observed in the data, compared to 70% for the Benchmark model. At the same time, the landlord credit model features a smaller rise in the homeownership rate, explaining only 6% of the rise in the data, compared to 35% for the Benchmark model. These results are consistent with the intuition in Figure 2 Panel (d) where landlord credit shifts the supply curve out.

Figure A.1: 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 4. 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.

The results holding parameters fixed would, however, make the model inconsistent with our empirical results.³⁸ To address this, we repeat the counterfactual that combines a credit expansion and decline in interest rates exercise in Section 4 to recalibrate $\sigma_{\omega,L}$ for the landlord credit model.³⁹ The resulting responses, denoted “Landlord Credit (Recalibrated),” follow a very similar pattern, explaining 70% of the observed rise in the price-rent ratio (as opposed to 70% without landlord credit), and 27% of the rise in the homeownership rate (as opposed to 35% in the baseline).


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 conservative and should provide a lower bound on the true contribution of credit over this period.

Robustness: No Landlord Credit Relaxation. Our main results with landlord credit relax *both* landlord and household credit in the boom, which is the most plausible assumption. In the less likely scenario in which landlord credit did not relax in the boom, our main landlord credit counterfactual would overstate the degree to which credit affects

³⁸We 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.

³⁹Roughly 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 the calibration subsection above for details). In recalibrating, we thus assume only half of the landlords get the subsidy.

Figure A.2: Robustness, No Relaxation of Landlord Credit Standards



Notes: Plots display perfect foresight paths following a relaxation of credit standards and a decline in interest rates. The “Landlord Credit (Recalibrated)” model applies the landlord credit extension recalibrating $\sigma_{\omega,L}$ under the new model, as in Figure A.1. The “Landlord Credit (HH Expansion Only)” uses the same model, but applies an alternative experiment in which household credit standards are relaxed and household mortgage interest rates decline, but landlord credit standards and mortgage interest rates are left unchanged.


house prices because supply would shift out less in the boom. In this section, we quantitatively evaluate this scenario to put a lower bound on the role of credit in the boom.

To check robustness to this concern, we compute the responses to an alternative boom-bust experiment in which we recalibrate the model as in our baseline landlord experiment, lower mortgage rates, and relax borrower credit standards, but do *not* relax landlord credit standards. The results of this experiment are displayed in Figure A.2. This figure shows that the alternative experiment delivers a path for the price-rent ratio that is slightly lower (explaining 60% vs. 70% of the observed increase), but a substantially larger increase in the homeownership rate (explaining 88% vs. 27% of the observed increase). The intuition for this finding is that the slope of our calibrated demand curve, estimated to match estimates from Berger et al. (2020), is substantially flatter than our calibrated supply curve. As a result, expansions of credit supply to borrowers, which shift the tenure demand curve, have a larger impact on homeownership than on the price-rent ratio. We conclude that our main results indicating strong effects of credit on house prices are robust, while our results on the impact of credit on homeownership may be understated if landlord credit standards were not relaxed during the boom period.

A.3 Extension: Flexible Saver Demand

In this section, we relax our assumption of fixed saver demand $H_{S,t} = \bar{H}_S$ and allow savers to freely trade housing, with $H_{S,t}$ as an additional control variable. This adds an

Figure A.3: 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 4. 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.

additional equilibrium condition from the saver first order condition

$$p_t^{\text{Saver}} = E_t \left\{ \Lambda_{t+1}^S \left[\underbrace{u_{h,t}^S / u_{c,t}^S}_{\text{housing services}} + \underbrace{(1 - \delta) p_{t+1}}_{\text{continuation value}} \right] \right\} \quad (\text{A.5})$$

This expression is nearly identical to the borrower’s condition (5) 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 A.3 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 4. 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 the expansion in demand due to credit. 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 empirical IRFs in Section 2.2, 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, returning the ratio of these responses to the values observed in our empirical estimates.⁴⁰ Even with a perfectly frictionless saver margin, the recalibrated saver model still explains 51% of the observed rise in the price-rent ratio from changes in the price and quantity of credit alone. While this response is about significantly smaller than the 70% 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.

A.4 Testing the Model Against Johnson (2020)

To further test the empirical performance of our model, we simulate a version of the main empirical exercise in Johnson (2020). Johnson (2020) exploits empirical variation driven by a divergence of PTI limit policy between Fannie Mae and Freddie Mac. Specifically, beginning in 1999, Freddie Mac appears to impose a PTI limit of 50 for a large share of borrowers, while Fannie Mae appears to leave PTI limits effectively unconstrained, with no clear bunching at any level. As a result, Johnson (2020) finds that the share of loans originated by Freddie Mac in a county in 1998, prior to the policy change, negatively predicts both the share of loans issued in that county with PTI ratios exceeding 50, and also predicts lower house price growth.

To replicate this experiment in our model, we consider a hypothetical house price that would hold in an area with lending standards designed to mimic those of Freddie Mac. To mimic Freddie Mac standards, as shown in Johnson (2020) Figure I(A), we assume

⁴⁰The 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.

that following the credit expansion, Freddie Mac imposes a PTI limit of 65% (as in our benchmark experiment) for half of borrowers, but imposes a lower PTI limit of 50% for the other half. This 50-50 split is chosen to visually match the evidence in Johnson (2020) Figure I(A), but the exact split is not particularly important, as we will be studying the ratio of the effect on house prices to the effect on the share of borrowers with PTI limits in excess of 50%. For example, while a smaller share with the 50% PTI limit should reduce the response of both variables, the impact on the ratio of the two should be second order.

The resulting “Freddie Mac” house price satisfies (5) using an alternative measure of $\mathcal{C}_{B,t}$ that takes into account the alternative PTI limit. Specifically:

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

where

$$\begin{aligned} \mathcal{C}_{B,t+1}^{\text{Freddie}} &= \mu_{B,t} F_t^{LTV, \text{Freddie}} \theta^{LTV} \\ F_t^{LTV, \text{Freddie}} &= \frac{1}{2} \left(\Gamma_e(\bar{e}_t) + \Gamma_e(\bar{e}_t^{50}) \right) \\ \bar{e}_t^{50} &= \frac{\theta^{LTV} p_t H_{B,t}^* (r_{B,t}^* + \nu + \alpha)}{(50\% - \omega) y_{B,t}}. \end{aligned}$$

We also compute the share of borrowers with PTI ratios in excess of 50 in the Benchmark and Freddie Mac economies:

$$\begin{aligned} \text{Share}_{>50} &= 1 - \Gamma_e(\bar{e}_t^{50}) \\ \text{Share}_{>50}^{\text{Freddie}} &= \frac{1}{2} \left(1 - \Gamma_e(\bar{e}_t^{50}) \right) = \frac{1}{2} \text{Share}_{>50}. \end{aligned}$$

With these variables defined in the model, we compare their response to the empirical estimates of Johnson (2020). Table III of Johnson (2020) displays the estimated coefficients of the change in the share of loans with PTI exceeding 50 (“Share DTI > 50”) after the policy change on the lagged Freddie Mac share, reporting point estimates of -2.90 with controls and -3.29 without controls. Table V reports the estimated coefficients of the percent change in house prices over the two quarters following the change in policy (Jun 1999 - Dec 1999) on the lagged Freddie Mac share, finding point estimates of -1.87 with controls or -2.59 without controls. Taking the ratio of these yields values of 0.645 with controls or 0.787 without controls.


Table A.1: Comparison: Model vs. Johnson (2020)

Experiment	Ratio
Johnson (2020), Controls	0.645
Johnson (2020), No Controls	0.787
Model, Credit Standards Only	0.557%
Model, Credit Standards + Rates	0.717%
Model, Full Boom	0.570%

Notes: top two rows display ratios of coefficients in Table V of Johnson (2020) to Table III of Johnson (2020). Bottom three rows display ratios of $(\log p_t - \log p_t^{\text{Freddie}})$ and $(\text{Share}_{>50} - \text{Share}_{>50}^{\text{Freddie}})$ under various experiments: “Credit Standards Only” (see Figure 7 Panel (a)), “Credit Standards + Rates” (see Figure 7 Panel (b)), and “Full Boom” (see Figure 8 Panel (a)).

To compute model equivalents, we the ratio of $(\log p_t - \log p_t^{\text{Freddie}})$ and $(\text{Share}_{>50} - \text{Share}_{>50}^{\text{Freddie}})$ under our various boom experiments. For the most direct comparison to Table V of Johnson (2020), we evaluate this ratio two quarters following the unexpected shock (change in policy), with the results displayed in Table A.1. In our experiment that only loosens credit standards (Figure 7 Panel (a)) — our most direct analogue to the policy change — we find a ratio of 0.541, which is very close to the main ratio 0.645 of Johnson (2020), and within the empirical confidence interval. Incorporating the decline in rates as in Figure 7 Panel (b) increases this ratio to 0.681, while moving to our “Full Boom” experiment of Figure 8 Panel (a) yields 0.518, showing that these results are robust to allowing for additional sources of variation.

Figure A.4: Comparison: Model vs. Johnson (2020)



Notes: Plots display perfect foresight paths. Results are summarized numerically at the 2Q horizon in Table A.1. Experiments map to previous figures as follows: “Credit Standards Only” (see Figure 7 Panel (a)), “Credit Standards + Rates” (see Figure 7 Panel (b)), and “Full Boom” (see Figure 8 Panel (a)). For data definitions see notes for Figure 1 and Table 2.

Figure A.4 displays our model ratios for additional horizons, as well as the numerator and denominator used in computing these ratios. The resulting paths show that our ratios are not highly sensitive to the choice of a two-quarter horizon. The figure plots the implied ratio over the first 14Q of each experiment, showing that the ratios remain close to that of Johnson (2020), particularly for our main model analogue, the “Credit Standards Only” experiment.

This multi-horizon plot can also be used to compare our model implications to an alternative set of house price regressions that Johnson (2020) estimates over a longer 14Q horizon. These regressions, shown in Johnson (2020) Table VI, find a coefficient on the lagged Freddie Mac share of -6.69 with controls and -8.86 without controls. Dividing by the coefficients found in the Share DTI > 50 regressions in Table III yield larger ratios of 2.31 with controls and 2.69 without controls. These ratios are not directly comparable to our results in Figure A.4 because the denominator in these empirical ratios is the *initial* response of the high-PTI share, while the numerator is the 14Q response of house prices. As shown above, house prices were trending upward in the data over this period. Since the share of borrowers with high PTI ratios should be increasing in the house price, the coefficient on the high-PTI share regressions should also have been increasing over this period. As a result, the ratio using 14Q changes in both numerator and denominator should be smaller, just as predicted by the model in Figure A.4.⁴¹ Still, to the extent that the long-run ratio may be higher, Figure A.4 shows that our results do not overstate these longer horizon ratios, implying that our results are if anything conservative in this case.

A.5 Robustness: Varying $\sigma_{\omega,L}$

Since our empirical results in Section 2 present our main findings in terms of the ratio of the price-rent response to the homeownership rate response, we provide a mapping between these ratios and our main model results. For this exercise, rather than jointly matching the entire path of price-rent ratio and homeownership rate responses to our LS instrument, as in Section 4.2, we instead calibrate $\sigma_{\omega,L}$ to match a specific ratio of the price-rent response to the homeownership rate response at a fixed horizon, which we choose to be two years.⁴² By varying the target ratio of the price-rent response to the homeownership rate response, we can provide a clear mapping between the various

⁴¹In principle, we could have taken the ratio of the 14Q difference in house prices to a shorter-horizon change in the high-PTI share. However, one shortcoming of our parsimonious model is that house prices jump on arrival of the policy rather than adjusting gradually, implying that our model would not be a good laboratory for measuring the size of this bias.

⁴²Unlike in our baseline estimation, we only re-estimate $\sigma_{\omega,L}$, while holding fixed our estimates of the persistence and size of the interest rate shock from our initial estimation.

Table A.2: Results, Boom Experiments, by Target Ratio

Experiment	Price-Rent	Homeown.	Loan-Inc.
Peak Data Increase	51.5%	3.3pp	72.2%
<i>Credit Relaxation (Share of Peak Data Increase)</i>			
Ratio = 1	14%	109%	40%
Ratio = 2	24%	59%	47%
Ratio = 3	28%	38%	49%
Ratio = 4	30%	27%	50%
Ratio = 5	31%	21%	51%
Ratio = 6	32%	17%	51%
Ratio = 7	33%	14%	52%
Ratio = 8	33%	12%	52%
Ratio = 9	33%	11%	52%
Ratio = 10	34%	10%	52%
<i>Credit Relaxation + Decline in Rates (Share of Peak Data Increase)</i>			
Ratio = 1	31%	201%	57%
Ratio = 2	52%	114%	71%
Ratio = 3	61%	75%	76%
Ratio = 4	65%	54%	79%
Ratio = 5	68%	42%	81%
Ratio = 6	70%	34%	82%
Ratio = 7	71%	29%	82%
Ratio = 8	72%	25%	83%
Ratio = 9	72%	22%	83%
Ratio = 10	73%	19%	84%

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 5. 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.

ratios we computed in our empirical analysis and the implied results fitted to that ratio.

The results of this exercise are displayed in Table A.2. Our baseline estimates, for which the implied ratio at the two-year horizon is close to six, are unsurprisingly very close to the “Ratio = 6” rows of the table. For robustness, recall that our empirical point estimates for this ratio were at least three across all specifications and horizons. Mapping this into the “Ratio = 3” rows of the table, we observe that these minimum estimates

would still deliver strong effects of credit on house prices, with a credit relaxation alone explaining 28% of the observed rise in price-rent ratios, and a combination of credit relaxation and low rates explaining 61% of this observed rise. Similarly, our bootstrapped lower bounds of our confidence intervals for this ratio were at least two across all specifications and horizons. Mapping this into the “Ratio = 2” rows, we find that a credit relaxation alone would explain 24% of the observed rise in price-rent ratios, while a combination of credit relaxation and low rates would explain 52% of this observed rise. At the same time, estimates using these lower ratios deliver larger responses of the homeownership rate to credit, in line with the intuition in Section 1.


We conclude that calibrating our model to match any of our empirical results, not only our baseline LS estimates, would lead to strong measured effects of credit on house prices that do not differ dramatically from our baseline estimates.

A.6 Additional Model Results


This section presents additional model results referenced in the main text. Figure A.5 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.5 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.6 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.6, where we see a

Figure A.5: Credit Relaxation by Borrower Heterogeneity



(a) Credit Relaxation Experiment




(b) Credit Relaxation + Rates Experiment

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.

larger response of both prices and rents in the model with fixed construction supply.

Figure A.6: Credit Relaxation by Construction Supply Elasticity



Notes: Figure shows responses of the indicated variables 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. Our data sources and construction are described in detail in the Data Availability Statement in our replication readme, with a shorter summary of things a reader would find useful to know here.

B.1.1 LS Instrument Data Set

We create an annual panel of CBSAs and metropolitan divisions, henceforth referred to as “CBSAs” from 1990 to 2017.⁴³ The main data set is comprised of 402 CBSAs. For most of our analysis in the main text we focus on data from 1995 to 2017 (the TW data is only available for all CBSAs in 1994 and we lag this variable by one year). We either take annual averages or choose quarter 2 as the observation for a given year, as appropriate (quarter 2 for HPI and QCEW employment, 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 and metropolitan divisions from 1976 to 2018; for cases where we need a house price index for a CBSA with metro divisions, we aggregate the metro division indices to create a CBSA-wide index. We use FHFA house price indices (all transactions for all CBSAs, purchase only for the largest 100 CBSAs, and expanded data for the 50 largest CBSAs) 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

⁴³CBSAs are collections of counties. For 11 of the CBSAs there are “metropolitan divisions” which are smaller subdivisions of the larger CBSA, such as Orange County and Los Angeles in the greater LA-Orange County CBSA or Dallas and Ft. Worth in the larger Dallas-Ft.Worth CBSA. There are 11 CBSAs with metropolitan divisions. Whenever possible we use metropolitan divisions and drop the larger CBSA, although in some cases some data (e.g., a homeownership rate) will come at the CBSA level for a CBSA with metropolitan divisions, in which case we use the CBSA.

we map to CBSAs. We are able to map to a quarterly panel for 53 CBSAs beginning in 1989 and 62 CBSAs beginning in 1994.

- Homeownership rates:
 - Our first 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 4% in either of the two closest censuses, we flag the series as bad and drop it from our main 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. We also do not include some locations in the final HVS analysis sample if (1) the geographies at which the HVS data are available do not match the TW data or (2) the HVS data does not have a complete panel from 1994-2017 for a geography.
 - Our second homeownership rate measure is our new microdata-based homeownership rate. Because the creation and benchmarking of this data series is involved, we cover it separately in Appendix C. This data series provides homeownership rates for a balanced panel of 390 CBSAs from 1994-2017.
- Credit Data: We use credit data from the Home Mortgage Disclosure Act microdata, which we collapse to the CBSA level to create the fraction of originations within 5% of the CLL. We use the same data restrictions as Loutskina and Strahan (2015) in creating this fraction: There has to be a positive and non-missing loan amount, a positive, non-missing, and non-top-coded applicant income, a non-missing state and county, be coded as a conventional loan (non-Veterans Administration, non-Federal Housing Administration, non-Farm Service Agency, and non-Rural Housing Ser-

vice), and finally be originated or denied. We have experimented with other data restrictions and find that the exact restriction used does not meaningfully impact the results, which is why we simply follow LS.

- 2000 Population and Housing units are obtained from NHGIS.
- Housing supply elasticity: We use data from Saiz (2010) which we crosswalk from his MSA definitions to our CBSAs using principal cities.
- Employment and industry shares: We use the quarterly 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. We then use quarter 2 as the annual observation.

The CBRE Torto-Wheaton rent index merits additional discussion. As mentioned in the main text, it measure 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.

From the merged data set, we create several samples. The two main samples are the “HVS Sample” and the “GG Microdata Sample,” which are used in the main text. The HVS Sample has 41 CBSAs from 1994-2017 (with results form 1995-2017 due to using lagged outcome variables as a control), while the GG Microdata Sample has 62 CBSAs

form 1994-2017. The GG Microdata sample includes all CBSAs for which we have TW rent data. The HVS sample drops CBSAs that (1) have a “bad” HVS series due to a significant CBSA definition change as detailed above, (2) that have an incomplete HVS panel (e.g. are not covered for all 24 years), and (3) where the HVS and TW data cover different geographies. The CBSAs in each sample are listed in Table B.1.

In this Appendix, we run a number of analyses on expanded samples which are described as we come to them. The largest sample includes 390 CBSAs for which we have house prices, GG homeownership, and other necessary data to run the full analysis with only house prices and GG homeownership rates.

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 original 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 and includes 370 CBSAs for the GG microdata-based homeownership rate and 47 when using the HVS homeownership rate.

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 this into the same data set as for the LS instrument. Our final data set includes 258 CBSAs from 1990 to 2017. We are missing one CBSA from the MS data set, Poughkeepsie NY, because it was absorbed into another CBSA using the 2013 CBSA definitions and thus does not match to one of the CBSAs in our analysis.

B.2 LS Instrument Details and Robustness


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

The Loutskina-Strahan instrument is 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. 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

Table B.1: CBSAs in Main Analysis Samples

CBSA Name	In HVS Sample	In GG Microdata Sample
Albuquerque NM Metropolitan Statistical Area	No	Yes
Anaheim-Santa Ana-Irvine CA Metropolitan Division	No	Yes
Atlanta-Sandy Springs-Roswell GA Metropolitan Statistical Area	Yes	Yes
Austin-Round Rock TX Metropolitan Statistical Area	Yes	Yes
Baltimore-Columbia-Towson MD Metropolitan Statistical Area	Yes	Yes
Birmingham-Hoover AL Metropolitan Statistical Area	Yes	Yes
Boston-Cambridge-Newton, MA-NH	No	Yes
Charlotte-Concord-Gastonia NC-SC Metropolitan Statistical Area	Yes	Yes
Chicago-Naperville-Elgin, IL-IN-WI	No	Yes
Cincinnati OH-KY-IN Metropolitan Statistical Area	Yes	Yes
Cleveland-Elyria OH Metropolitan Statistical Area	Yes	Yes
Columbus OH Metropolitan Statistical Area	Yes	Yes
Dallas-Plano-Irving TX Metropolitan Division	No	Yes
Denver-Aurora-Lakewood CO Metropolitan Statistical Area	Yes	Yes
Detroit-Warren-Dearborn, MI	Yes	Yes
El Paso TX Metropolitan Statistical Area	No	Yes
Fort Lauderdale-Pompano Beach-Deerfield Beach FL Metropolitan Division	No	Yes
Fort Worth-Arlington TX Metropolitan Division	No	Yes
Greensboro-High Point NC Metropolitan Statistical Area	Yes	Yes
Greenville-Anderson-Mauldin SC Metropolitan Statistical Area	No	Yes
Hartford-West Hartford-East Hartford CT Metropolitan Statistical Area	Yes	Yes
Houston-The Woodlands-Sugar Land TX Metropolitan Statistical Area	Yes	Yes
Indianapolis-Carmel-Anderson IN Metropolitan Statistical Area	Yes	Yes
Jacksonville FL Metropolitan Statistical Area	Yes	Yes
Kansas City MO-KS Metropolitan Statistical Area	Yes	Yes
Las Vegas-Henderson-Paradise NV Metropolitan Statistical Area	Yes	Yes
Los Angeles-Long Beach-Glendale CA Metropolitan Division	No	Yes
Louisville/Jefferson County KY-IN Metropolitan Statistical Area	Yes	Yes
Memphis TN-MS-AR Metropolitan Statistical Area	Yes	Yes
Miami-Miami Beach-Kendall FL Metropolitan Division	No	Yes
Minneapolis-St. Paul-Bloomington MN-WI Metropolitan Statistical Area	Yes	Yes
Nashville-Davidson-Murfreesboro-Franklin TN Metropolitan Statistical Area	Yes	Yes
Nassau County-Suffolk County NY Metropolitan Division	No	Yes
Newark NJ-PA Metropolitan Division	No	Yes
New York-Jersey City-White Plains NY-NJ Metropolitan Division	No	Yes
Oakland-Hayward-Berkeley CA Metropolitan Division	No	Yes
Oklahoma City OK Metropolitan Statistical Area	Yes	Yes
Orlando-Kissimmee-Sanford FL Metropolitan Statistical Area	Yes	Yes
Oxnard-Thousand Oaks-Ventura CA Metropolitan Statistical Area	No	Yes
Philadelphia PA Metropolitan Division	No	Yes
Phoenix-Mesa-Scottsdale AZ Metropolitan Statistical Area	Yes	Yes
Pittsburgh PA Metropolitan Statistical Area	Yes	Yes
Portland-Vancouver-Hillsboro OR-WA Metropolitan Statistical Area	Yes	Yes
Providence-Warwick RI-MA Metropolitan Statistical Area	Yes	Yes
Raleigh NC Metropolitan Statistical Area	No	Yes
Richmond VA Metropolitan Statistical Area	Yes	Yes
Riverside-San Bernardino-Ontario CA Metropolitan Statistical Area	Yes	Yes
Sacramento-Roseville-Arden-Arcade CA Metropolitan Statistical Area	Yes	Yes
St. Louis MO-IL Metropolitan Statistical Area	Yes	Yes
Salt Lake City UT Metropolitan Statistical Area	No	Yes
San Antonio-New Braunfels TX Metropolitan Statistical Area	Yes	Yes
San Diego-Carlsbad CA Metropolitan Statistical Area	Yes	Yes
San Francisco-Redwood City-South San Francisco CA Metropolitan Division	No	Yes
San Jose-Sunnyvale-Santa Clara CA Metropolitan Statistical Area	Yes	Yes
Seattle-Tacoma-Bellevue, WA	Yes	Yes
Tampa-St. Petersburg-Clearwater FL Metropolitan Statistical Area	Yes	Yes
Tucson AZ Metropolitan Statistical Area	Yes	Yes
Tulsa OK Metropolitan Statistical Area	Yes	Yes
Urban Honolulu HI Metropolitan Statistical Area	Yes	Yes
Virginia Beach-Norfolk-Newport News VA-NC Metropolitan Statistical Area	Yes	Yes
Washington-Arlington-Alexandria, DC-VA-MD-WV	Yes	Yes
West Palm Beach-Boca Raton-Delray Beach FL Metropolitan Division	No	Yes

Figure B.1: Loutskina-Strahan Instrument LP Impulse Response For House Prices and Rents



Notes: 95% confidence interval shown in bars. The figure shows panel local projection estimates of the response of the indicated outcomes for the indicated samples to the LS instrument $ShareNearCLL_{i,t} \times \%ChangeInCLL_t$ as estimated using equation (1). Control variables include $ShareNearCLL_{i,t}$ and its lag and lags of the instrument and outcome variables, and regressions are weighted by 2000 population. Standard errors are clustered by CBSA.


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.

Figure B.1 shows the impulse response of rents and prices separately for various sample (the HVS sample of 41 CBSAs, the GG sample of all 62 CBSAs with TW data, and, for prices, the full sample of 390 CBSAs). One can see that essentially all of the IRF to our credit shock comes from prices: rents respond by a statistically insignificant amount. Furthermore, the three price samples have similar IRFs.

The limited response of rents motivates using the full sample of 390 CBSAs with prices and the GG microdata-based homeownership rather than price to rent and homeownership to see if the results are different for the expanded sample. This is shown in Figure B.2. We can see that the results are quite similar to the baseline results in Figure 3 for years 0-2 but have a slightly larger homeownership response and a smaller price response in years 3-4. The ratio (not the inverse ratio) is generally larger than our baseline estimates for the GG microdata, ranging from 28 to infinity in years 0-2 and 6.6 to 7.8 in years 3-4 when the price response is smaller and homeownership response is larger.

Figure B.3 shows the response of the HVS homeownership rate for various samples and shows that the results are similar across samples. In particular, it show the baseline (41 CBSAs), a version that does not condition on having rents from the TW index in the

Figure B.2: Loutskina-Strahan Instrument LP Impulse Responses: Expanded Sample House Prices Only (No Rents)



Notes: 95% confidence interval shown in bars. The figure shows panel local projection estimates of the response of the indicated outcomes to the LS instrument $ShareNearCLL_{i,t} \times \%ChangeInCLL_t$ as estimated using equation (1). Control variables include $ShareNearCLL_{i,t}$ and its lag and lags of the instrument and outcome variables, and regressions are weighted by 2000 population. Panel (a) shows the price and homeownership rate for the GG homeownership rate with standard errors clustered by CBSA. Panel (b) shows the inverse ratio $\beta_k^{PRR} / \beta_k^{HOR}$, with standard errors block bootstrapped by CBSA.


sample but continues to drop CBSAs with a large change in the HVS homeownership rate (54 CBSAs), and finally a version that does not require a full balanced panel (67 CBSAs, unbalanced from 1992-2017). The results are similar across the samples.

Figure B.4 repeats the main analysis in Figure 3 but includes time-varying controls for employment and industry shares from the QCEW. For the GG microdata-based homeownership rate, this uses log total employment and two-digit industry shares. Because of the smaller sample we cannot control for two-digit industry shares in the HVS and instead use log total employment and one-digit industry shares. For the HVS, the impulse response for both PRR and HOR are moderately smaller, leading to larger ratios in all periods as the denominator effects is stronger than the numerator. The results for the GG homeownership rate sample are similar, with both impulse responses slightly smaller and a larger ratio. This implies that time-varying city characteristics related to employment and industry composition are not driving the results.

Figure B.5 adds the employment and industry controls to the 390-CBSA specification that uses house prices and GG homeownership as in Figure B.2. Again, the results are similar and the ratios are slightly higher, from 35 to infinity in periods 0-2 and 9.1 to 11.0 in periods 3-4.

Figure B.6 compares the impulse responses for price for various price indices. In the main text we use the CoreLogic price index. The figure shows this alongside the responses

Figure B.3: Loutskina-Strahan Instrument LP Impulse Responses: HVS Various Samples



(a) Point Estimates

Notes: 95% confidence interval shown in bars. The figure shows panel local projection estimates of the response of the HVS homeownership rate to the LS instrument $ShareNearCLL_{i,t} \times \%ChangeInCLL_t$ as estimated using equation (1), for the indicated samples. Control variables include $ShareNearCLL_{i,t}$ and its lag and lags of the instrument and outcome variables, and regressions are weighted by 2000 population with standard errors clustered by CBSA. The “baseline” sample is the main sample of HVS CBSAs with good homeownership rates and TW rents. The “All CBSAs in HVS” sample is all HVS CBSAs with good homeownership rates but a full balanced panel. The “All CBSAs in HVS, Unbalanced” is the same as “All CBSAs in HVS” but drops the requirement that there be a full balanced panel.


of the FHFA all transactions (which includes sales and appraisals) for 388 CBSAs and the FHFA purchase only index for 99 CBSAs. The figure shows that our results are robust to the price index used.

Finally, Figure B.7 repeats the analysis for 1991-2017 instead of 1995-2017 and drops the requirement of a fully balanced panel, although we do require each CBSA in the sample to have at least 20 years of data. Again, the results are similar to the baseline analysis in Figure 3, suggesting that starting the analysis in 1995 to have a full balanced panel of TW rents is not an important sample restriction.

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 for growth variables. We run the same regression as DK, then transform the impulse response from log changes to log levels by cumulating

Figure B.4: Loutskina-Strahan Instrument LP Impulse Responses: Employment and Industry Share Controls




Notes: 95% confidence interval shown in bars. The figure shows panel local projection estimates of the response of the indicated outcomes to the LS instrument $ShareNearCLL_{i,t} \times \%ChangeInCLL_t$ as estimated using equation (1). Control variables include $ShareNearCLL_{i,t}$ and its lag, lags of the instrument and outcome variables, log employment, and industry employment shares (one-digit for HVS and two-digit for GG), and regressions are weighted by 2000 population. Panels (a) and (b) show the price/rent and homeownership rate for the HVS and GG homeownership rates, respectively, with standard errors clustered by CBSA. Panels (c) and (d) show the inverse ratio $\beta_k^{PRR} / \beta_k^{HOR}$ for the HVS and GG homeownership rates, respectively, with standard errors block bootstrapped by CBSA.

Figure B.5: Loutskina-Strahan Instrument LP Impulse Responses: Expanded Sample House Prices Only (No Rents) With Employment and Industry Share Controls




Notes: 95% confidence interval shown in bars. The figure shows panel local projection estimates of the response of the indicated outcomes to the LS instrument $ShareNearCLL_{i,t} \times \%ChangeInCLL_t$ as estimated using equation (1). Control variables include $ShareNearCLL_{i,t}$ and its lag, lags of the instrument and outcome variables, log employment, and two-digit industry employment shares, and regressions are weighted by 2000 population. Panel (a) shows the price and homeownership rate for the GG homeownership rate with standard errors clustered by CBSA. Panel (b) shows the inverse ratio $\beta_k^{PRR} / \beta_k^{HOR}$, with standard errors block bootstrapped by CBSA.

Figure B.6: Loutskina-Strahan Instrument: FHFA House Price Index




Notes: 95% confidence interval shown in bars. The figure shows panel local projection estimates of the response of the indicated price indices to the LS instrument $ShareNearCLL_{i,t} \times \%ChangeInCLL_t$ as estimated using equation (1). Control variables include $ShareNearCLL_{i,t}$ and its lag and lags of the instrument and outcome variables. Regressions are weighted by 2000 population and standard errors block bootstrapped by CBSA.

Figure B.7: Loutskina-Strahan Instrument LP Impulse Responses: 1991-2017 Unbalanced Panel



Notes: 95% confidence interval shown in bars. The figure shows panel local projection estimates of the response of the indicated outcomes to the LS instrument $ShareNearCLL_{i,t} \times \%ChangeInCLL_t$ as estimated using equation (1). Control variables include $ShareNearCLL_{i,t}$ and its lag and lags of the instrument and outcome variables, and regressions are weighted by 2000 population. Panels (a) and (b) show the price/rent and homeownership rate for the HVS and GG homeownership rates, respectively, with standard errors clustered by CBSA. Panels (c) and (d) show the inverse ratio $\beta_k^{PRR} / \beta_k^{HOR}$ for the HVS and GG homeownership rates, respectively, with standard errors block bootstrapped by CBSA. This figure differs from the main text because it includes data from 1991 onwards instead of 1995 onwards and is an unbalanced panel, although we require each CBSA in the sample to have at least 20 years of data.

Figure B.8: Di Maggio-Kermani APL Preemption Reduced Form: Replicating DK's Specification



Notes: 95% confidence interval shown in bars. The figure shows estimates of the estimates of β_k for each indicated year and outcome variable (price or the GG microdata-based homeownership rate) estimated from equation (2), with the instrument being $Z_i = APL_{2004} \times OCC_{2003}$. 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. All regressions are weighted by 2000 population and standard errors are clustered by CBSA, as in the original DK paper.

the coefficients from the log changes regression.


B.3.1 Replicating DK's Exact Specification

To compare our results directly to Figure 3 of DK, Figure B.8 replicates DK's exact specification as best we can by showing the impulse response in growth rates rather than levels. The blue circles show the CBSA-level data, while the green triangles show the analysis using DK's full county-level data set. One can see that the CBSA and County level data are similar, and our estimates are close to DK's published results. For the remainder of this appendix, we return to using cumulated impulse responses.

B.3.2 Robustness

Figure B.9 shows the same analysis as in the main text, that is using the GG microdata-based homeownership rate with a cumulated IRF, but adding in controls for log employment and two-digit industry shares from the QCEW. The results are very similar to Figure 4 in the main text; the house price response and the homeownership rates response are both statistically and economically insignificantly smaller. Because of the small magnitude of the homeownership rate response, the ratio is at 1

Figure B.9: Di Maggio-Kermani APL Preemption Reduced Form With Employment and Industry Share Controls



Notes: 95% confidence interval shown in bars. The figure shows estimates of the cumulative sum from 2003 of β_k for each indicated year and outcome variable (price or the GG microdata-based homeownership rate) estimated from equation (2), with the instrument being $Z_i = APL_{2004} \times OCC_{2003}$. 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, the fraction of HUD-regulated lenders interacted with a dummy for APLs, the log of total employment from the QCEW, and 2-digit employment shares from the QCEW. All regressions are weighted by 2000 population and standard errors are clustered by CBSA, as in the original DK paper.


Figure B.10 shows the Di Maggio and Kermani analysis using the HVS homeownership rate sample. This yields 53 CBSAs instead of 370. The HVS yields a much lower ratio (between 1.7 and 3.0), but has extremely wide confidence intervals, to the point that one cannot reject a ratio of infinity. We thus do not make much of the lower ratios for the HVS data.

B.4 MS Instrument Details and Robustness

Figure B.11 reestimates our main MS regression controlling for log of total employment and two-digit industry shares from the QCEW. The impulse response for prices peaks lower, but the qualitative pattern of a large response of prices and a small response of homeownership still holds. The ratio is between 22 and 23 in 2004 and 2005.


Figure B.12 reestimates our main MS regression using homeownership rate data from the HVS. This yields 36 CBSAs instead of 258. As mentioned in the main text, using the HVS gives smaller estimates between 1.7 and 2.4 in 2004 and 2005, but with very wide confidence intervals, to the point that one cannot reject a ratio of infinity. We thus do not make much of the lower ratios for the HVS data.

Figure B.10: Di Maggio-Kermani APL Preemption Reduced Form With HVS Homeownership Rate




Notes: 95% confidence interval shown in bars. The figure shows estimates of the cumulative sum from 2003 of β_k for each indicated year and outcome variable (price or the GG microdata-based homeownership rate) estimated from equation (2), with the instrument being $Z_i = APL_{2004} \times OCC_{2003}$. 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. All regressions are weighted by 2000 population and standard errors are clustered by CBSA, as in the original DK paper.

Figure B.11: Mian-Sufi PLS Expansion Reduced Form With Employment and Industry Share Controls



Notes: 95% confidence interval shown in bars. The figure shows estimates of the effect of a city's NCL share on the indicated outcome (price or the GG microdata-based homeownership rate) based on estimating equation (2) with the instrument being $Z_i = NCLShare_i^{2002}$ and 2002 being the base year. The regressions control for the log of total employment from the QCEW, and 2-digit employment shares from the QCEW. All standard errors are clustered by CBSA and all regressions are weighted by housing units as in the original MS paper.

Figure B.12: Mian-Sufi PLS Expansion Reduced Form With HVS Homeownership Sample



Notes: 95% confidence interval shown in bars. The figure shows estimates of the effect of a city's NCL share on the indicated outcome (price or the GG microdata-based homeownership rate) based on estimating equation (2) with the instrument being $Z_i = NCLShare_i^{2002}$ and 2002 being the base year. All standard errors are clustered by CBSA and all regressions are weighted by housing units as in the original MS paper.

B.5 National Price-Rent Ratio Construction

This appendix describes the data construction for Figure 1. 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.

C A New Measure of the Homeownership Rate

C.1 Data Sources

Our data construction relies on two sources.

Infutor. For information on the inhabitant of a property, we use Infutor’s Total Consumer ID Plus (CRD4) data set. These data contain information on the address history of the majority of adults in the United States. The data trace individual address histories for up to 10 addresses or 30 years, whichever is shorter. For each historical address and each individual, Infutor provides the first and last name, the first date at which the individual lived at that address, and data on the address itself. We define the end date of a residential spell as the next start date that is strictly greater than the current start date, or as January 1st, 2020 (beyond the sample end date) for the final residence for that individual.

ZTRAX. For information on the owner of a property and its characteristics, we rely on Zillow’s Transaction and Assessment Database (ZTRAX). These data provide information on the buyers and sellers on deeds transactions, including first and last name, the property address, the transaction type, and the sale date. The data also provide information on the owner on dates when the property is assessed for taxes, including the address, assessment date, and owner name. We also get property characteristics from the assessment.

Public Housing. We obtain data on all public housing units from the department of Housing and Urban Development. Our data was obtained on September 24th, 2020 from: <https://hudgis-hud.opendata.arcgis.com/datasets/HUD::public-housing-buildings/about>

C.2 Data Preparation

Before performing our main data merge, we take several steps to prepare and clean the data. We first describe the generic steps for name separation, name cleaning, and address standardization that we apply to all data sets, before listing additional cleaning steps dataset-by-dataset.

Name Separation. A major challenge is that some data, particularly the deeds and assessor data, combine multiple individuals in a single name, or put both first and last name in the last name field. As a result, the “last name” variable often includes both first and

last names of multiple individuals. In addition, the data include several inconsistent ways of recording multiple names varying with whether the first name comes before the last name, whether the names are separated with a comma, and whether individuals with the same last name are grouped with a single last name and multiple first names. To address these issues, we run perform the following steps to separate names into different individuals and into first and last names in all of our data sets that include name data:

1. For observations that include a non-missing first name, we leave the data as is, since these observations usually represent a single individual, correctly formatted.
2. We classify observations that appear to be some type of corporate entity, and also leave these observations as is. We assign names to this category if they include an exact match of any of the following words:

AND, ASS, ASSET, ASSS, BARCLAYS, BEAR STEARNS, BUILT, CHURCH, CITY, CO, CU, DEUTSCHE, DEV, ENTER, ENTS, FIN, GA, HOMES, HUD, IGLESIA, JP MORGAN, LEASE, LL, LOAN, LP, MONEY, MOR, MORGAN STANLEY, NB, OWNER, PROP, PROPS, RE, REAL, REO, RLTY, SONS, STORE, STRUCTURES, TAX, TITLE, UNION, US, USA, VA

or if the string ends with an exact match of any of the following words:

ADVISORS, ALTISOURCE, AMERICA, AMERICAN, ASSETS, ASSN, ASSOC*, BANK, BANKING, BAPTIST, BK, BKNG, BLDG, BLDR, BLDRS, BORROWER, BROADCASTING, BUILDERS, CAPITAL, CENTER, CITIGROUP, CITIMORTGAGE, CNTY, COALITION, COMMUNIT*, CONGREGATION, CONSORTIUM, CONST, CONSTR, CONSTRUCTION, CONTRACTOR, CONTRACTORS, CORP, CORPORATION, COUNTRYWIDE, COUNTY, CREDIT, DEPOSIT, DEPT, DEVEL, DEVELOP*, ENTERPRISES, ENTPR, EQUITIES, FARGO, FB, FBS, FCU, FDIC, FED, FEDERAL, FINANCE, FINANCIAL, FINL, FIRST, FNB, FNDG, FNMA, FSB, FSB, FUND, FUNDING, GMAC, GRP, HOLDING, HOLDINGS, HOME, HOMEOWNERS, HOUSING, HSBC, INC, INDEPENDENT, INTL, INVEST, INVESTMENT, INVESTMENTS, INVESTORS, INVST, INVTRS, JLC, JPMORGAN, LDNG, LENDER, LENDERS, LENDING, LIABILITY, LIMITED, LLC, LLP, LNDNG, LOANS, LTD, MANAGEMENT, METHODIST, MGMT, MINISTRIES, MORT, MORTG, MORTGAGE, MSAC, MTG*, MUTUAL, NATIONAL, NATIONSTAR, NATL, OPENDOOR, PARTNERS, PARTNERSHIP, PROPERTIES, PROPERTY, PURCHASE, REALTY, RELOCATION, RENEWAL, RENTAL, RENTALS, RESIDENTIAL, SB, SEC, SECRETARY, SECURITIES, SERIES, SERVICES, SERVICING, SFR, SOLUTIONS, SRVC, SUNTRUST, SVCNG, SVCS, UNDERWRITING, UNITED, VENTURES, VETERANS, WAREHOUSE.

The “*” symbol above is a wildcard including any additional letters that finish the word, as in typical regular expressions. For any entry without that wildcard, we require an exact match from the start to the end of an individual word in the string.

3. For remaining observations that include a comma, we base our name separation approach on the pattern of commas and ampersands in the last name.
 - (a) For strings with a single name followed by a comma and then a set of additional names separated by ampersands, we assume that this string represents a single last name followed by a set of first names. For each first name we create a new observation with that first name and the common last name. For example, the string "SMITH, JOHN & JANE" would be mapped into two (first name, last name) pairs: (JOHN, SMITH), (JANE, SMITH).
 - (b) For strings with one or more patterns of two names separated by commas, each of which are separated from each other by ampersands, we assume that each "name, name" pair represents a last name followed by a first name. We create a new observation with that first and last name for each "last, first" pair. For example, the string "SMITH, JOHN & JONES, JANE" would be mapped into two (first name, last name) pairs: (JOHN, SMITH), (JANE, JONES).
 - (c) For strings that first feature a comma, followed by some alternative comma and ampersand pattern that does not fall into one of these two cases (e.g., "SMITH, JOHN, JANE"), we treat the initial name before the comma as the common last name, and create separate observations using each name following the initial comma as the first name. For example, the string "SMITH, JOHN, JANE" would be mapped into two (first name, last name) pairs: (JOHN, SMITH), (JANE, SMITH).
4. For remaining observations that do not include a comma, it is very difficult to distinguish first and last names. Because of this, we create a new observation for each word in the original last name string, and set the last name equal to this word. This means that our deeds data assign to the last name variable to some names that are actually first names. However, we note that because the Infutor data is almost always correctly formatted into first and last names, there is little danger of accidentally matching a first name in the deeds data to a first name in the Infutor data, so this type of error should have little impact on our measure of owner occupancy.

Name Cleaning. Once combined names have been separated into individual (first name, last name) observations, we clean the name data. We first make adjustments for prefixes and suffixes. For some observations there can be inconsistencies in how common prefixes such as "Le" or "La" are treated. A common discrepancy is that one data set will include a space between the prefix and the remainder of the name, while the other will not. To

address this, for a set of common name prefixes (DE, ST, MC, VAN, LE, LA, DEL) we remove any space between the prefix and the remainder of the last name.

Similarly, there can also be inconsistencies in whether and how suffixes such as “Jr.” are recorded. To address this, we remove any suffix in the set (JR, SR, II, III, IV) that form an exact match at the end of the last name string.

Last, we first convert all names to uppercase, to address possible inconsistencies in how names are cased. We then remove any non-alphanumeric characters other than dashes, whitespace, or ampersands, and strip any whitespace at the start or end of the string.

Adjusting Residential Spells. In our Infutor historical address data, we often observe that individuals appear to repeatedly move in and out of the same address. Since this may reflect measurement error rather than actual moves, we redefine an individual’s spell at a given address as a continuous stretch between the first time they begin residence at an address to the last time they end residence at that address.

Address Standardization. As a final update to all data sets, we standardize all addresses so that they will match more consistently in our merge. We use a standardization service from SmartyStreets to obtain standardized versions of each address. We associate with each standardized address a unique address ID number that we will use in our merge. Since properties can be owned either at the individual unit (e.g., apartment) level or at the building level, we create standardized versions each address after removing the portions of the address relating to the subunit, which we denote the *building address*. For example, if “123 Main Street Apt 1” is the full address of the unit, we would define the building address as “123 Main Street.” We add unique address ID numbers to each building address not already present elsewhere in our data.

Address History Data. In addition to the steps listed above, we take the following steps to clean and prepare our Infutor address history dataset.

1. We construct the end date for each residential spell as the start date of the next residential spell that is strictly after the current one.
2. Using the start and end dates defined in the previous step, we often observe that individuals appear to repeatedly move in and out of the same address. Since this may reflect measurement error rather than actual moves, we redefine an individual’s spell at a given address as a continuous stretch between the first time they begin residence at an address to the last time they end residence at that address.

3. Properties may be an entire building or a subunit (e.g., apartment) of that building. It may be possible that an inhabitant lives in a subunit, while the owner owns the entire building. To deal with this issue, we define for each property the building ID as described above.

Deeds Data. This section describes how we combine our original deeds, assessor, and public housing data to create a final set of deeds transactions.

1. We drop all deeds transactions that do not have broad document type “D” or “H,” which cover sales transactions. This drops transactions with broad document type “M” or “F,” which relate to mortgage and foreclosure-related transactions, respectively.
2. We set the end date of each ownership spell as the next recorded transaction date. For the final transaction recorded for a given property, we set the end date to a date beyond the end of our sample.
3. While we generally use the most recent buyers to identify the names of the current owners, this approach alone would not be able to define owner names prior to the first recorded transaction. To obtain the initial owners prior to the first recorded transaction, we create an additional deeds record for the property with date prior to the start of our sample (e.g., January 1, 1900) that lists the *sellers* on the first recorded transaction in our data as the *buyers* on this newly created observation. This ensures that the initial sellers will be recorded as the owners of the property for all periods prior to the first recorded transaction.
4. We create additional deeds records corresponding to the recorded owners in our assessment data. For each assessment, we create a new deed using the recorded owner from the assessment as the buyer, with a purchase date equal to the previous recorded purchase in our transaction data (or January 1, 1900 if missing), and a sale date equal to the next recorded purchase in our transaction data (or January 1, 2100 if missing). These extra assessment-based records are particularly useful for cases where we have no transaction data at all, such as the case of a very longstanding owner who has not transacted for decades.
5. We create additional deeds records for all addresses in our public housing data that will identify these properties as owned by a government entity. For these records we use a placeholder for the last name that will never match with any last name in the address history data, and hence will never show up as an owner-occupied

property. However, by assigning an owner to these properties, we will correctly be able to identify residential spells in these properties as non-owner-occupied.

C.3 Data Merge

After constructing our two data sets containing information on residential spells and ownership, we merge the two data sets. We perform the merge by address ID, keeping all residents and all owners who have ever been associated with a given address, as well as the start and end dates of both the residential spells and of homeownership. These merged data are thus indexed at the (address, inhabitant, owner) level. This is a many-to-many nature merge, so that all inhabitants are matched with all possible owners.

Our initial merge will not capture properties that are owned at the building level but inhabited at the subunit (e.g., apartment) level, because the address ID of the building will differ from the address ID of the subunit. To deal with this, we separately merge our deeds data with our residential spells data using address ID as the merge key for the deeds data, and *building address ID* as the merge key for the residential spells data, where building address is defined as above. We append this second merged data set to our initial data set after dropping any observations where the building address ID and address ID are the same (in which case the observation is already in the data set).

Our initial merge includes many observations that are not relevant because they include entirely non-overlapping dates, either because the inhabitant moved into the property after the owner sold it, or the inhabitant left the property before the owner bought it. Because these observations cannot possibly influence either the numerator (whether the inhabitant and owner have the same last name at a given date) or the denominator (whether we have data on both an inhabitant and an owner at a given date), we drop these observations.

For each remaining (address, inhabitant, owner) observation, we check whether we obtain an exact match between the last name of the inhabitant and the last name of the owner. To address the possibility that first and last names were reversed, we also check whether we obtain an exact match between first name of the inhabitant and last name of the owner *and* between last name of the inhabitant and first name of the owner. In either of these two cases, we say that this property owner is an *owner-occupier* and that the property is owner occupied throughout that owner's ownership spell. We chose this date convention over an alternative date convention where we only define a property as owner occupied during the overlapping period between the ownership spell and the inhabitant's residential spell because from our hand checking of the data we believe that

the dates associated with the deeds data are much more precise than the dates in the residential history data. However, we acknowledge that this approach will misclassify events where an owner spends part of their ownership spell inhabiting a property and part of their ownership spell renting it.

C.4 Homeownership Rate Calculation

Given our owner occupied flag for each (address, inhabitant, owner) observation, we can aggregate to obtain a geographic time series of homeownership rates as follows.

1. Fix a date, denoted DATE, at which we are going to evaluate the homeownership rate.
2. Find all observations where DATE is weakly between the date at which the owner purchased the property and sold the property.
3. For each remaining address, compute an occupancy flag for whether the address has at least one registered inhabitant and at least one registered owner.
4. For each remaining address, compute an owner-occupancy flag as the maximum over the values of the owner-occupancy flag over all (address, inhabitant, owner) observations. This determines whether we identify a property as owner-occupied at a given date.
5. At the geographic level, sum over values of the occupancy flag and the owner-occupancy flag. Divide the sum of the owner-occupancy flag (number of owner-occupied units) by the sum of the occupancy flag (number of units for which we have both resident and owner information) to obtain an estimate of the geographic homeownership rate on date DATE.
6. Repeat for DATE equal to each date of interest.

In our implementation, we computed the homeownership rate at the county (FIPS) level on the first day of each quarter from 1980:Q1 to 2019:Q4. We sum our totals of the owner-occupied and in-sample flags over the four quarters of each calendar year to obtain annual ratios, and drop any (county, year) observation with fewer than 25 in-sample properties. We will denote this initial measure for county i and year t as $HOR_{i,t}^{GG,raw}$.

C.5 Trend Adjustment

Our data construction in the previous sections provides a raw measure of the homeownership rate, corresponding to the share of units with non-missing owners and occupants that are owner-occupied. However, our data coverage changes over time, mostly due to the Infutor data increasing coverage and scope. This may create low frequency trends that do not match actual homeownership changes. Moreover, because of differences in coverage over time across counties, these time trends may vary from county to county, and will not be completely removed by a combination of date and county fixed effects

In this section, we describe how we use high-quality homeownership data available at low frequencies from the Decennial Census and the American Community Survey (ACS) to remove the low-frequency trend in raw homeownership rate data.

To motivate our procedure, our main approach to this trend adjustment is to update the low-frequency trend in our data to match a trend line that interpolates linearly over the 10 year periods between each Decennial Census. However, at the time our data were constructed, county-level homeownership rates in the 2020 Decennial Census were not available. To address this, we instead use the low-frequency trend in the ACS from 2005 onward to construct corrected trends over the end of our sample.

Our procedure for doing so is as follows.

1. To address any bias in the aggregate homeownership rate for each date, we remove time effects from our GG homeownership measure, and will replace these constants later on with the correct national homeownership rate at each date. Specifically, we remove time effects from our homeownership measures using

$$\widetilde{HOR}_{i,t}^s = HOR_{i,t}^s - \overline{HOR}_t^s,$$

where s represents the source of the homeownership rate, which is either our newly constructed GG measure (“GG, raw”), the ACS, or the Decennial Census, where the latter two will be used to adjust the low frequency trends in our data later on. The time averages \overline{HOR}_t^s are computed as weighted averages of $HOR_{i,t}^s$ across counties i using constant county weights equal to the number of occupied units in that county in the 2005 ACS.

2. Compute the county-level trend in the homeownership measure in each county over the period where the ACS is available using the regression:

$$\widetilde{HOR}_{i,t}^s = \hat{\alpha}_i^s + \hat{\beta}_i^s \times t + \varepsilon_{i,t},$$

where we run a separate regression for each county i , and the source s is one of our GG measure, the ACS, or the HVS. If at least three observations are available for county i , we compute the trend homeownership rate over the ACS sample as:

$$\widehat{HOR}_{i,t}^{s, \text{post-2005 trend}} = \hat{\alpha}_i^s + \hat{\beta}_i^s \times t.$$

3. Construct our measure of the correct low-frequency trend in the data, denoted $\widehat{HOR}_{i,t}^*$ as follows:

(a) For years 2005 or later, compute:

$$\widehat{HOR}_{i,t}^* = \widehat{HOR}_{i,t}^{ACS} + \gamma_i = \hat{\alpha}_i^s + \hat{\beta}_i^s \times t + \gamma_i,$$

where the addition of the constant:

$$\gamma_i = \widetilde{HOR}_{i,2010}^{Census} - \widehat{HOR}_{i,2010}^{ACS},$$

ensures that our trend line is exactly equal to the Decennial Census measure $\widetilde{HOR}_{i,2010}^{Census}$ in 2010.

(b) For years 2000 or earlier, we linearly interpolate between homeownership rates in the time-demeaned Decennial Census $\widetilde{HOR}_{i,t}^{Census}$. For example,

$$\widehat{HOR}_{i,1993}^* = 0.7 \times \widetilde{HOR}_{i,1990}^{Census} + 0.3 \times \widetilde{HOR}_{i,2000}^{Census}.$$

(c) For the years 2001-2004, we linearly interpolate between $\widehat{HOR}_{i,2000}^*$ and $\widehat{HOR}_{i,2005}^*$, where these two values are computed using the steps above (and $\widehat{HOR}_{i,2000}^* = \widetilde{HOR}_{i,2000}^{Census}$).

4. Repeat the procedure in the previous step to construct the low-frequency trend with the same structure in our “GG, raw” data, $\widehat{HOR}_{i,t}^{GG, \text{raw}}$. The idea is that we will remove this low-frequency trend from $\widetilde{HOR}_{i,t}^{GG, \text{raw}}$ and replace it with $\widehat{HOR}_{i,t}^*$. To be precise, we compute $\widehat{HOR}_{i,t}^{GG, \text{raw}}$ as follows:

(a) For years 2005 or later, compute the fitted value from a linear time trend:

$$\widehat{HOR}_{i,t}^{GG, \text{raw}} = \hat{\alpha}_i^{GG} + \hat{\beta}_i^{GG} \times t.$$

- (b) For years 2000 or earlier, we linearly interpolate between homeownership rates $\widetilde{HOR}_{i,d}^{GG,raw}$, where $d \in \{1980, 1990, 2000\}$ is a Decennial Census year. For example,

$$\widehat{HOR}_{i,1993}^{GG,raw} = 0.7 \times \widetilde{HOR}_{i,1990}^{GG,raw} + 0.3 \times \widetilde{HOR}_{i,2000}^{GG,raw} .$$

- (c) For the years 2001-2004, we linearly interpolate between $\widehat{HOR}_{i,2000}^{GG,raw}$ and $\widehat{HOR}_{i,2005}^{GG,raw}$, where these two values are computed using the steps above (and $\widehat{HOR}_{i,2000}^{GG,raw} = \widetilde{HOR}_{i,2000}^{GG,raw}$).

5. Compute a trend-adjusted measure of the homeownership rate as:

$$HOR_{i,t}^{GG,*} = \underbrace{\overline{HOR}_t^{Agg}}_{\text{national HOR}} + \underbrace{\widetilde{HOR}_{i,t}^{GG,raw}}_{\text{raw data}} - \underbrace{\widehat{HOR}_{i,t}^{GG,raw}}_{\text{raw data trend}} + \underbrace{\widehat{HOR}_{i,t}^*}_{\text{Census/ACS trend}} \quad (C.1)$$

where \overline{HOR}_t^{Agg} is the national homeownership rate, obtained from the US Census Bureau (FRED code RHORUSQ156N). To understand these expressions, note that e.g., (C.1) begins with the national homeownership rate, which since our remaining series are all demeaned will ensure that our data will always aggregate to the correct national number when using 2005 ACS occupied units as weights. Next, we add the demeaned data series $\widetilde{HOR}_{i,t}^{GG,raw}$, and subtract off the county-specific low frequency trend in our “GG, raw” data $\widehat{HOR}_{i,t}^{GG,raw}$. Last, we add back in the “correct” low-frequency trend from the Decennial Census and ACS $\widehat{HOR}_{i,t}^*$.

6. To aggregate from the county level to the CBSA level, we compute weighted averages of our $HOR_{i,t}^{GG,*}$ measures for each CBSA j , using the 2013 mappings from counties to CBSA (source: NHGIS), and using the population of each county, interpolated between Decennial Census years, as the weight.

The results of this procedure is the the CBSA-level series $HOR_{j,t}^{GG,*}$.

C.6 Data Validation

After constructing our homeownership rate series $HOR_{j,t}^{GG,*}$, we next seek to validate it by comparing it to the American Community Survey, a high quality data set but one available only since 2005, and to the Housing Vacancy Survey, the best existing public series containing years prior to 2005. To remove mechanical sources of common variation due to

Table C.1: Homeownership Rate Comparison

Statistic	ACS	HVS	GG
Standard deviation (levels)	0.70%	1.87%	0.33%
Standard deviation (1Y differences)	0.94%	2.18%	0.31%
Autocorrelation (levels)	0.089	0.331	0.692
Autocorrelation (1Y differences)	-0.395	-0.108	0.686
Standard deviation (deviation from ACS)	–	2.36%	0.94%
Correlation with ACS (levels)	1.000	0.058	0.261
Correlation with ACS (1Y differences)	1.000	0.027	0.121
Correlation with ACS (3Y differences)	1.000	0.036	0.296
Correlation with ACS (5Y differences)	1.000	0.081	0.380

Notes: This table presents summary statistics from homeownership series from the American Community Survey (ACS), Housing and Vacancy Survey (HVS), and our newly constructed GG-Microdata series for the 2015-2017 period in which all three series overlap. Statistics are equally weighted across CBSAs and time. The sample includes all CBSAs with at least ten years of data for the ACS, HVS, and GG series.

our trend adjustment procedure, we remove year and geographic fixed effects (weighted by ACS occupied units in 2005), as well as a linear time trend for each CBSA. This has the additional benefit of isolating the variation that would be left over in regression analyses, which often remove geographic and time fixed effects and time trends. We compare these homeownership rate series on the overlapping sample for which all three homeownership rates are available for at least ten years. This sample contains 960 observations from 75 CBSAs over the period 2005 to 2017.⁴⁴

To compare these series, we present summary statistics in Table C.1, and a full set of CBSA-level comparisons for each CBSA in our overlapping ACS/HVS/GG sample in Figures C.1 through C.6. Each figure presents the three versions of the homeownership rate for a particular CBSA. Since we remove a geographic effect and a linear time trend, all homeownership rates by construction have mean zero and no linear trend. However, we observe that the remaining variation, which is likely the most important for empirical analyses with fixed effects, varies widely across measures.

The top panel of Table C.1 displays statistics for the individual series. The HVS measure is by far the most volatile series, exhibiting twice the volatility of the ACS series in both levels and differences. This is not surprising as the HVS is built off of a supplement to Current Population Survey, which samples roughly 72,000 units. In 2021 there were roughly 142 million units in the US housing stock, meaning the HVS is built off of

⁴⁴There are 89 CBSAs in the HVS. Dropping CBSAs with major redefinitions leads to 80. Dropping CBSAs with under 10 years of data leads to 75.

a roughly 0.05% sample. Visual analysis of the figures by CBSA reveal reveals that this volatility is due to large swings that appear mostly uncorrelated with either the ACS or GG-Microdata homeownership rate. This series also displays negative autocorrelation in first differences, which is consistent with the presence of measurement error.

By contrast, the GG-Microdata series is the least volatile of all the series, with less than half the volatility of the ACS series in both levels and differences. This is likely due to the fact that our series is not a randomly resampled draw of households, but includes all of our data at each date. In contrast, while the ACS uses a larger sample than the HVS, it still only represents a random subsample of individuals far smaller than the actual population (around 2 million households, or a roughly 1.5% sample), which may incur sampling variation. Our series also displays the highest persistence in levels, and unlike the other series, does not display negative autocorrelation in first differences, providing further evidence that it has less measurement error.

The bottom panel of Table C.1 compares each series to the ACS. We observe that deviations between our GG series and the ACS are less than half of that for the HVS series. Similarly, the correlations between our series and the ACS in both levels and first differences are several times larger than for the HVS series. We conclude that our GG-Microdata series provides a much closer match to the ACS data over the overlapping sample when both are available. Looking at correlations with 3-year and 5-year differences, we observe that our GG-Microdata series exhibits large and growing correlations with changes in the ACS as the horizon becomes longer. This provides reassurance that the low volatility of our series is due to dampening noise and does not stem from a failure to capture true variation in the homeownership rate. We can also observe this visually from cases where the ACS homeownership rate exhibits large changes over the sample, such as Las Vegas-Henderson-Paradise, NV, Phoenix-Mesa-Scottsdale, AZ, or Salt Lake City, UT. These figures show that GG-Microdata series' reduction in noise does not come at the cost of understating actual movements in the ACS homeownership rate. Indeed, the GG-Microdata series generally tracks the ACS series very well, including in cases where the ACS series is far from stagnant.

Last, we provide evidence that, to the extent that our GG-Microdata and ACS series differ, the GG-Microdata may be the more accurate series. To this end, we regress:

$$ACS_{j,t+1} = \beta_0 + \beta_1 ACS_{j,t} + \beta_2 HVS_{j,t} + \beta_3 GG_{j,t} + \varepsilon_{j,t+1} \quad (C.2)$$

where $ACS_{j,t}$, $HVS_{j,t}$, and $GG_{j,t}$ are the homeownership rates in CBSA j at time t from the ACS, HVS, and GG-Microdata series, respectively. The results, displayed in Table C.2,

Table C.2: Regression Results


	$ACS_{j,t}$	$HVS_{j,t}$	$GG_{j,t}$
$ACS_{j,t+1}$	0.023 (0.042)	-0.000 (0.014)	0.596*** (0.079)
N	885		
Adjusted R^2	0.066		

Notes: This table presents results from an equal-weighted OLS regression of (C.2). The sample includes all CBSAs with at least ten years of data for the ACS, HVS, and GG series. Heteroskedasticity-robust standard errors are reported in parentheses.

show that our GG-Microdata series at time t is by far the strongest predictor of the ACS series at time $t + 1$ in the same CBSA, driving out all predictive power of the ACS series at time t itself. This implies that our GG-Microdata series is faithfully capturing the true “signal” in the homeownership rate, without the additional noise created by the ACS sampling scheme. We believe our method provides even more error reduction at the county level, where the random sampling of the ACS poses an even larger issue.

To summarize, on the overlapping sample for which the ACS, HVS, and GG-Microdata series are available, we find the GG-Microdata series to be the least volatile, most persistent, and most predictive of the next value of the ACS homeownership rate. All of these findings are consistent with the GG-Microdata series being a high-signal, low-noise measure of homeownership that is much more accurate than the HVS series, and may even improve on the accuracy of the ACS. These benefits should be considered alongside GG-Microdata’s expanded coverage that provides many more CBSAs than the HVS (390 vs. 75 with a continuous sample of more than three years) and a longer sample than the ACS, which begins only in 2005.

Figure C.1: Homeownership Rate Comparison by CBSA, Page 1




Notes: The series “ACS,” “GG-Microdata,” and “HVS,” correspond to homeownership rates measured using the American Community Survey, our new microdata-based measure, and the Census Housing and Vacancy Survey. For each series, we have removed CBSA and year fixed effects, as well as a linear time trend. The sample spans 2005-2017.

Figure C.2: Homeownership Rate Comparison by CBSA, Page 2




Notes: The series “ACS,” “GG-Microdata,” and “HVS,” correspond to homeownership rates measured using the American Community Survey, our new microdata-based measure, and the Census Housing and Vacancy Survey. For each series, we have removed CBSA and year fixed effects, as well as a linear time trend. The sample spans 2005-2017.

Figure C.3: Homeownership Rate Comparison by CBSA, Page 3




Notes: The series “ACS,” “GG-Microdata,” and “HVS,” correspond to homeownership rates measured using the American Community Survey, our new microdata-based measure, and the Census Housing and Vacancy Survey. For each series, we have removed CBSA and year fixed effects, as well as a linear time trend. The sample spans 2005-2017.

Figure C.4: Homeownership Rate Comparison by CBSA, Page 4




Notes: The series “ACS,” “GG-Microdata,” and “HVS,” correspond to homeownership rates measured using the American Community Survey, our new microdata-based measure, and the Census Housing and Vacancy Survey. For each series, we have removed CBSA and year fixed effects, as well as a linear time trend. The sample spans 2005-2017.

Figure C.5: Homeownership Rate Comparison by CBSA, Page 5



Notes: The series “ACS,” “GG-Microdata,” and “HVS,” correspond to homeownership rates measured using the American Community Survey, our new microdata-based measure, and the Census Housing and Vacancy Survey. For each series, we have removed CBSA and year fixed effects, as well as a linear time trend. The sample spans 2005-2017.

Figure C.6: Homeownership Rate Comparison by CBSA, Page 6



Notes: The series “ACS,” “GG-Microdata,” and “HVS,” correspond to homeownership rates measured using the American Community Survey, our new microdata-based measure, and the Census Housing and Vacancy Survey. For each series, we have removed CBSA and year fixed effects, as well as a linear time trend. The sample spans 2005-2017.