Supplemental Appendix: Do Credit Conditions Move House Prices?

Daniel L. Greenwald and Adam Guren*

March 2025

Abstract

This Supplemental Appendix is organized as follows. Appendix A contains additional results, details, derivations, and extensions for our structural model. Appendix B similarly provides additional results and detail for our cross-sectional empirical results. Last, Appendix C provides a full description of our construction of our GG-Microdata homeownership rate series, and presents its properties compared to existing alternatives.

*Greenwald: NYU Stern School of Business and NBER, 44 W 4th St #9-97, New York, NY 10012, dlg340@stern.nyu.edu. Guren: Boston University and NBER, 270 Bay State Road Room 406, Boston, MA 02215, guren@bu.edu.

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

$$p_t = \frac{\tilde{p}_t}{1 - \mathcal{C}_{B,t}} \tag{A.1}$$

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 $\mathcal{C}_{B,t}$ is in turn defined by

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

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:

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

In the second stage, all borrower-inhabited housing is divided into equal portions, including borrower-owned housing. Borrowers then draw their owner surplus shocks $\omega^B_{i,t}$. 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:

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

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

$$(h_{B,t}): \qquad q_t = (u^h_{B,t}/u^c_{B,t}), \tag{A.3}$$

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:

$$(M^*_{B,t}): \qquad 1 = \Omega^B_{M,t} + r^*_{j,t}\Omega^B_{X,t} + \mu_{B,t}, \tag{A.4}$$

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 $\Omega^B_{M,t}$ and of interest payments $\Omega^B_{X,t}$ satisfy:

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

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:

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

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

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

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:

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

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:

$$M^*_{L,t} \leq \theta^{LTV}_L p_t H^*_{L,t}.$$

The landlord's indifference condition for the marginal property now becomes

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

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

$$(M^*_{L,t}): \qquad 1 = \Omega^L_{M,t} + r^*_{j,t}\Omega^L_{X,t} + \mu_{L,t}.$$

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

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

symmetric to the borrower case.

The saver's budget constraint becomes:

$$\begin{aligned}
c_{S,t} \leq\ & \underbrace{(1-\tau)y_{S,t}}_{\text{after-tax income}} - \underbrace{p_t\left(H^*_{S,t} - H_{S,t-1}\right)}_{\text{net housing purchases}} - \underbrace{\delta p_t H_{S,t-1}}_{\text{maintenance}} + \underbrace{T_{S,t}}_{\text{rebates}} \\
& + \sum_{j\in\{B,L\}}\left\{\underbrace{\pi^{-1}(\bar{r}_j + \nu_j)M_{j,t-1}}_{\text{total payment}} - \underbrace{\rho_{j,t}\left(\exp(s_j\Delta_t)M^*_{j,t} - \pi^{-1}(1-\nu_j)M_{j,t-1}\right)}_{\text{net mortgage iss.}}\right\}
\end{aligned}$$

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 IV 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.

The results holding parameters fixed would, however, make the model inconsistent with our empirical results.[1] To address this, we repeat the counterfactual that combines a credit expansion and decline in interest rates exercise in Section IV to recalibrate $\sigma_{\omega,L}$ for the landlord credit model.[2] 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

[1] We 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.

[2] Roughly half of rental units are located in multifamily buildings too large to be affected by the changes in the CLL on which the LS instrument is based (see the calibration subsection above for details). In recalibrating, we thus assume only half of the landlords get the subsidy.

Figure A.1: Credit Standards + Falling Rates Experiment, Landlord Credit Extension

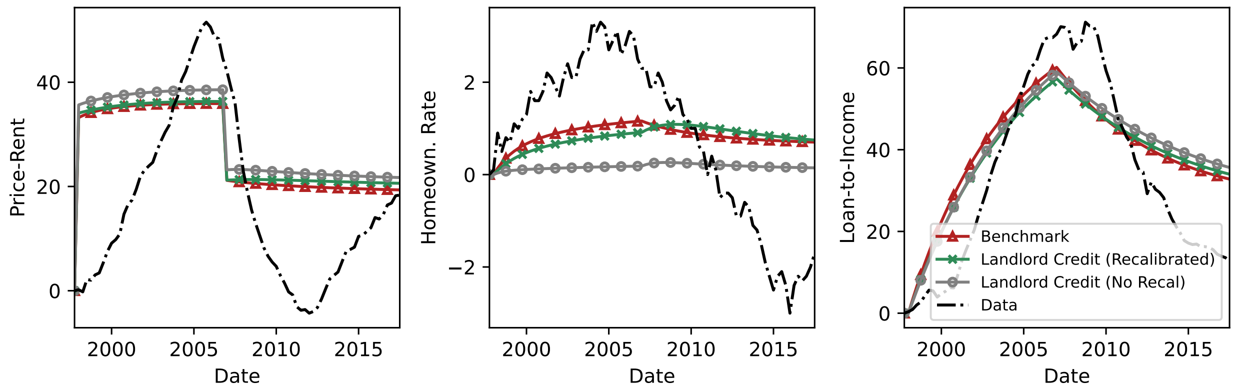

Notes: Plots display perfect foresight paths following a relaxation of credit standards and a decline in interest rates. The "Benchmark" model sets a value of $\sigma_{\omega,L}$ calibrated to match our empirical IRFs as in Section IV. The "Landlord Credit (No Recal)" model applies the landlord credit 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.

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 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 51% vs. 70% of the observed increase), but a substantially larger increase in the homeownership rate (explaining 120% vs. 27% of the observed increase). The intuition for this finding is that the slope of our calibrated demand curve, estimated to match estimates

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

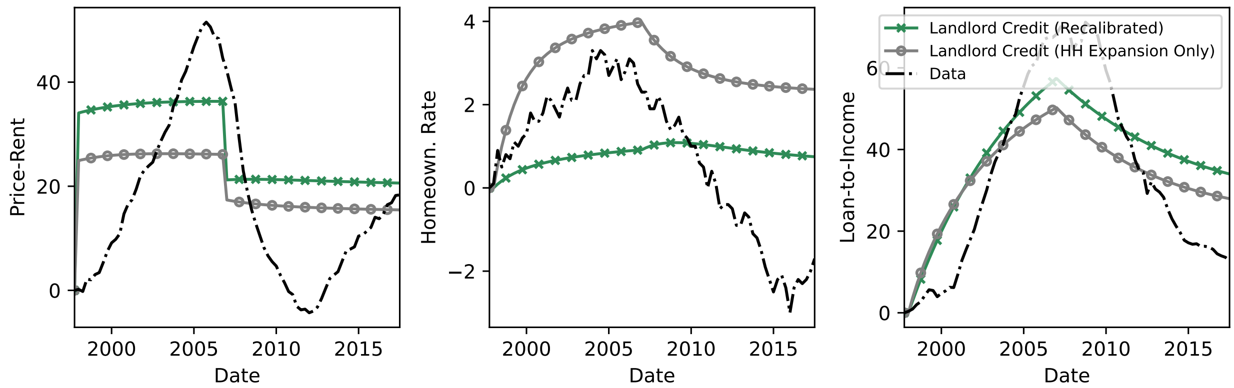

Notes: Plots display perfect foresight paths following a relaxation of credit standards and a decline in interest rates. The "Landlord Credit (Recalibrated)" model applies the landlord credit extension recalibrating $\sigma_{\omega,L}$ under the new model, as in Figure A.1. The "Landlord Credit (HH Expansion 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.

from Berger, Turner and Zwick (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 additional equilibrium condition from the saver first order condition

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

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 $p^{Saver} = p^{Demand} = p^{Supply}$ 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

Figure A.3: Credit Standards + Falling Rates Experiment, Saver Demand Extension

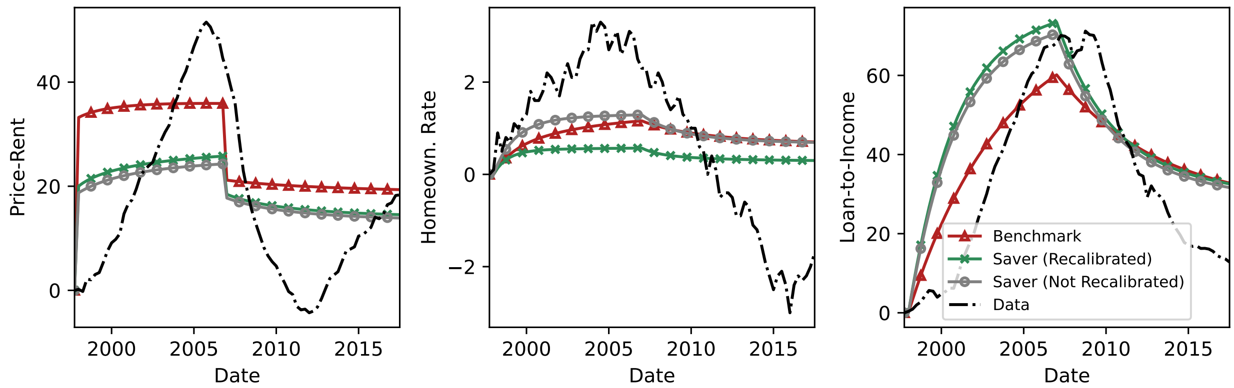

Notes: Plots display perfect foresight paths following a relaxation of credit standards and a decline in interest rates. The "Benchmark" model sets a value of $\sigma_{\omega,L}$ calibrated to match our empirical IRFs as in Section IV. The "Saver (Not Recalibrated)" model applies the flexible saver demand 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.

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 IV. 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 II.B, 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.[3] Even with a perfectly frictionless saver margin, the recalibrated saver model still explains 50% 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.

[3]The 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.

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, Piazzesi and Schneider (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 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:

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

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	0.567%
Model, Credit Standards + Rates	0.737%
Model, Full Boom	0.584%

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)).

where

$$\mathcal{C}_{B,t+1}^{\text{Freddie}} = \mu_{B,t} F_t^{LTV,\text{Freddie}} \theta^{LTV}$$

$$F_t^{LTV,\text{Freddie}} = \frac{1}{2}\left(\Gamma_e(\bar{e}_t) + \Gamma_e(\bar{e}_t^{50})\right)$$

$$\bar{e}_t^{50} = \frac{\theta^{LTV} p_t H_{B,t}^*(r_{B,t}^* + \nu + \alpha)}{(50\% - \omega) y_{B,t}}.$$

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

$$\text{Share}_{>50} = 1 - \Gamma_e(\bar{e}^{50})_t$$

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

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.

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)

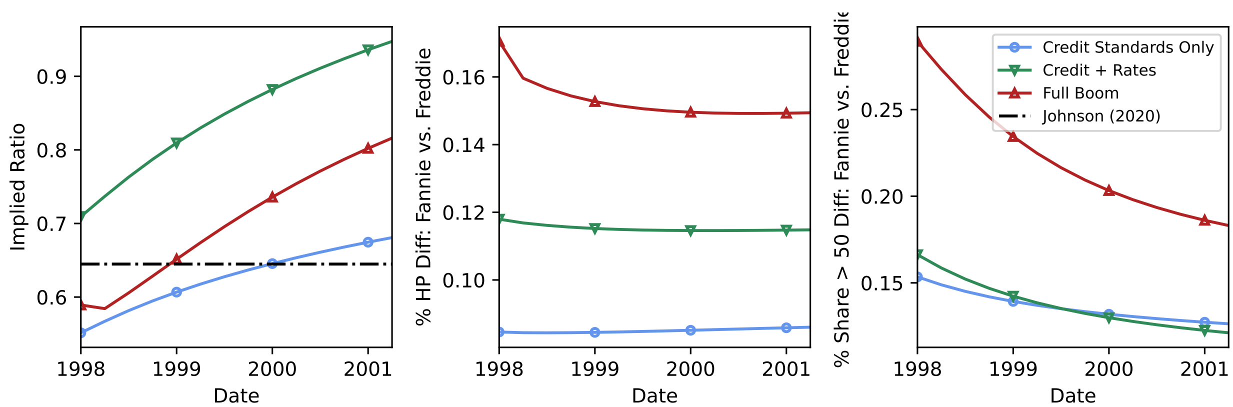

Notes: Plots display perfect foresight paths. Results are summarized numerically at the 2Q horizon in Table A.1. Experiments map to previous figures as follows: "Credit Standards Only" (see Figure 7 Panel (a)), "Credit Standards + Rates" (see Figure 7 Panel (b)), and "Full Boom" (see Figure 8 Panel (a)). 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.[4] 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 II 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 IV.B, 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.[5] By varying the target ratio of the price-rent response to the homeownership rate response, we can provide a clear mapping between the various 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

[4] In 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.

[5] Unlike 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%**	**3.3 pp**	**71.2%**
Credit Relaxation (Share of Peak Data Increase)			
Ratio = 1	14%	109%	41%
Ratio = 2	24%	59%	47%
Ratio = 3	28%	38%	50%
Ratio = 4	30%	27%	51%
Ratio = 5	31%	21%	52%
Ratio = 6	32%	17%	52%
Ratio = 7	33%	14%	53%
Ratio = 8	33%	12%	53%
Ratio = 9	33%	11%	53%
Ratio = 10	34%	10%	53%
Credit Relaxation + Decline in Rates (Share of Peak Data Increase)			
Ratio = 1	31%	202%	58%
Ratio = 2	52%	115%	72%
Ratio = 3	61%	75%	78%
Ratio = 4	66%	54%	81%
Ratio = 5	68%	42%	83%
Ratio = 6	70%	34%	84%
Ratio = 7	71%	29%	84%
Ratio = 8	72%	25%	85%
Ratio = 9	72%	22%	85%
Ratio = 10	73%	19%	85%

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 V. 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.

lower ratios deliver larger responses of the homeownership rate to credit, in line with the intuition in Section I.

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 "Credit Relaxation" experiment (Figure 7a in the main text) varying the borrower heterogeneity parameter $\sigma_{\omega,B}$. 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.

Figure A.5: Credit Relaxation by Borrower Heterogeneity

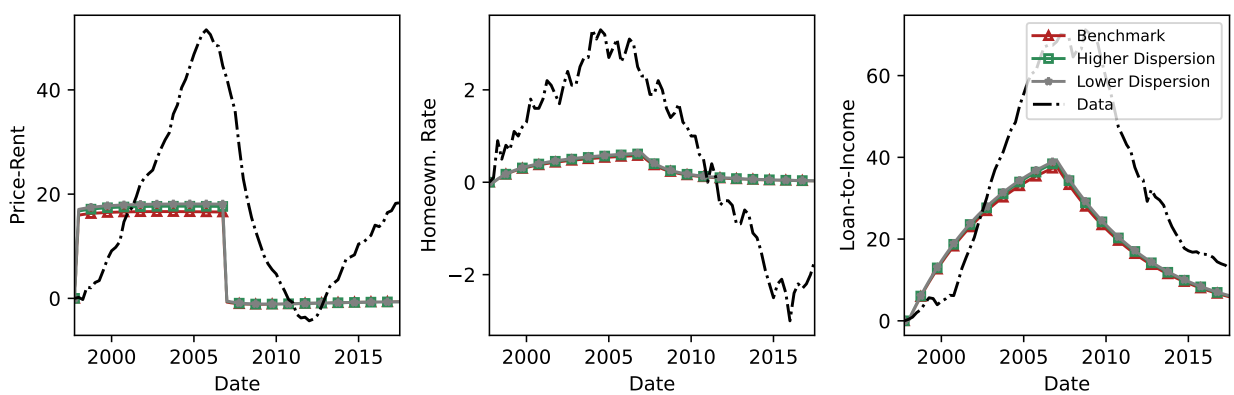

Notes: the figure shows responses to our Credit Relaxation experiment from Figure 7a varying the level of borrower dispersion (σ_B). The "Higher Dispersion" series sets $\sigma_{\omega,B}$ so that half as many renters ($0.64\%/2 = 0.32\%$) switch to ownership under the First Time Homeownership Subsidy, 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.

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

Figure A.6: Credit Relaxation by Construction Supply Elasticity

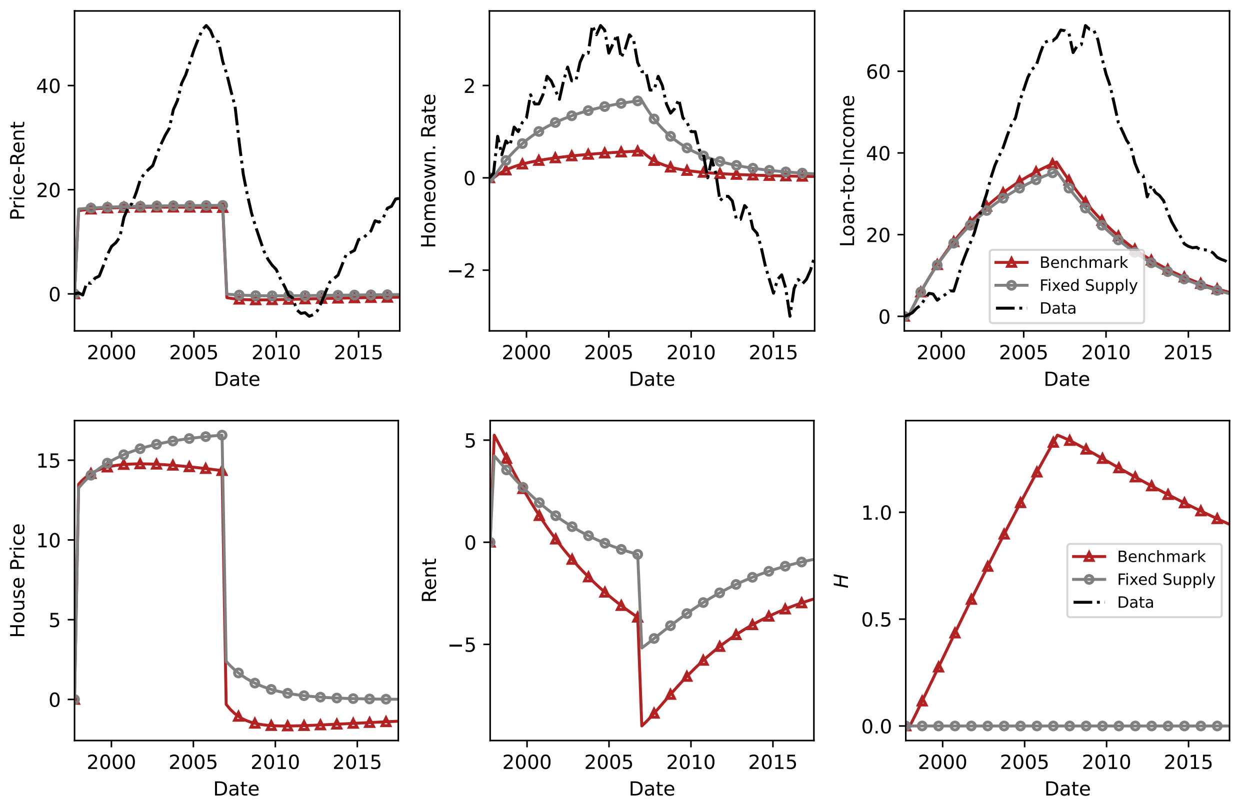

Notes: Figure shows responses of the indicated variables to our Credit Relaxation experiment (as also shown in Figure 7a) comparing the Benchmark model to a model with no construction sector and a fixed housing supply $\bar{H}$. For data definitions see notes for Figure 1 and Table 2. Additional definitions 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.[6] The main data set is comrpised 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 on 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 if 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 transaction 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 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.

[6]CBSAs 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.

- 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 Service), 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.

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

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 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

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

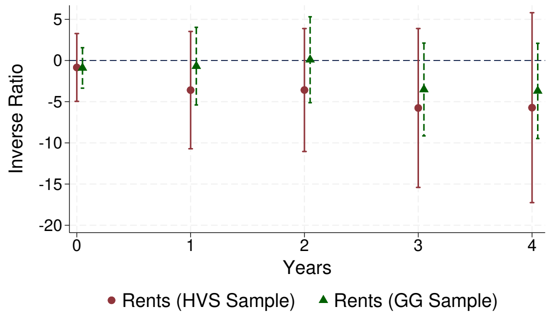

(a) Rents

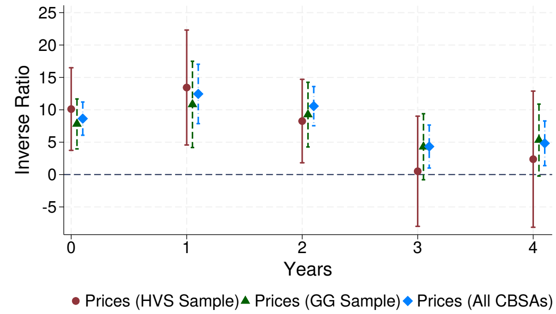

(b) Prices

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.

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 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

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

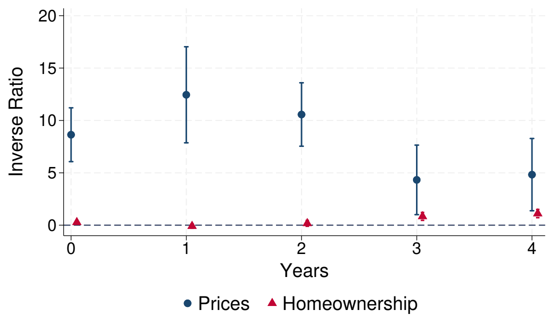

(a) Point Estimate

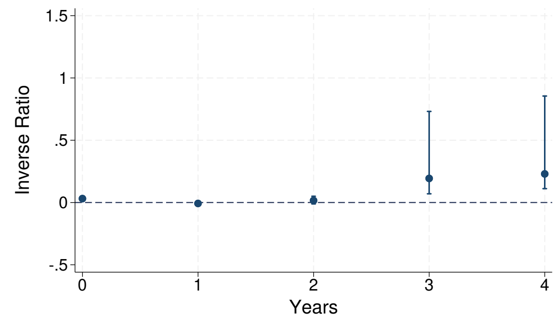

(b) Inverse Ratio

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.

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

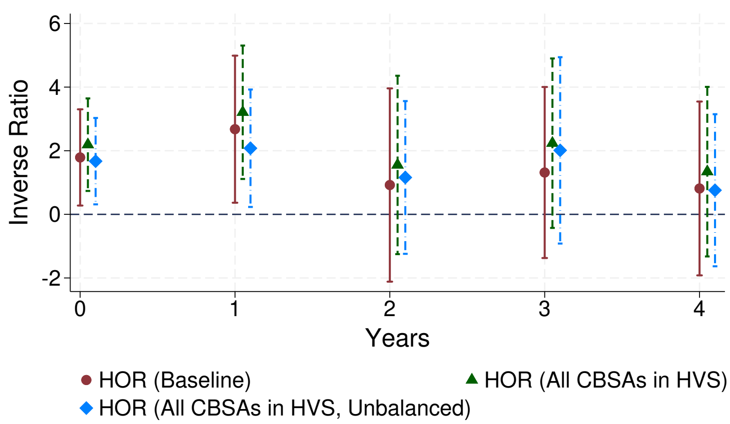

(a) Point Estimates)

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 "Al CBSAs in HVS" but drops the requirement that there be a full balanced panel.

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

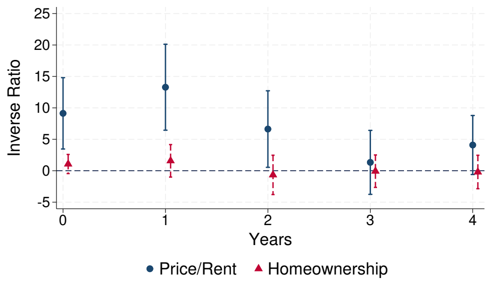

(a) Point Estimates (HVS)

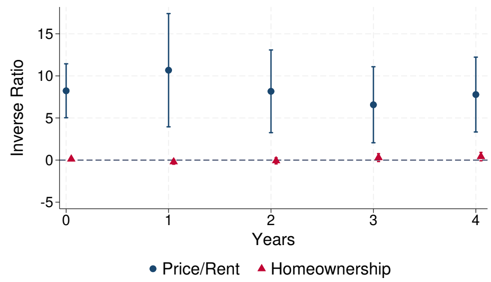

(b) Point Estimates (GG Microdata)

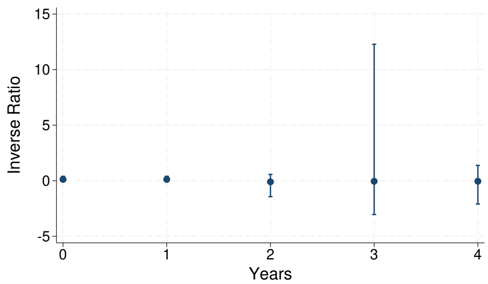

(c) Inverse Ratio (HVS)

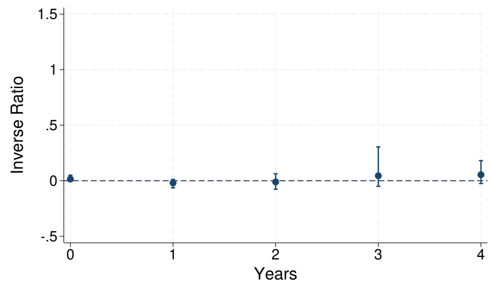

(d) Inverse Ratio (GG Microdata)

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.

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

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

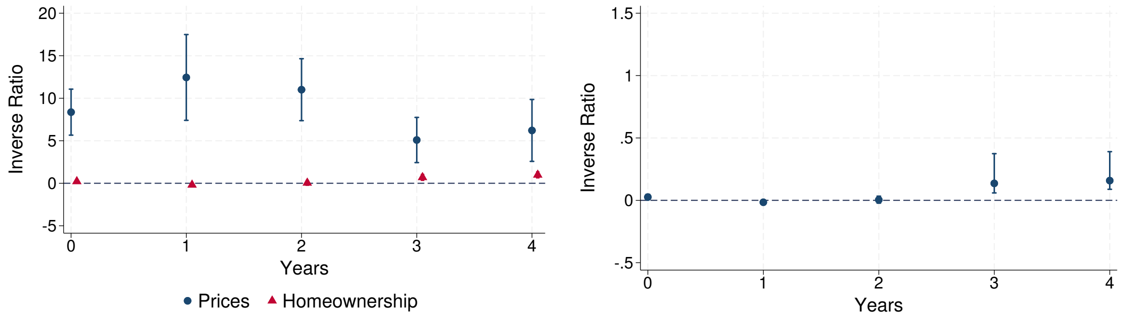

(a) Point Estimate

(b) Inverse Ratio)

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.

text we use the CoreLogic price index. The figure shows this alongside the responses 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 the coefficients from the log changes regression.

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

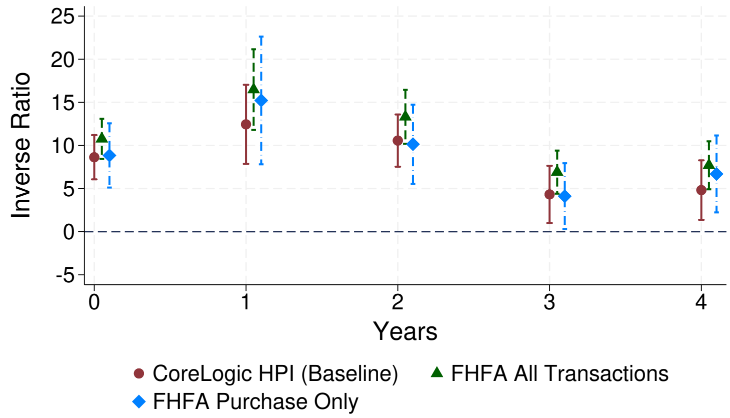

Notes: 95% confidence interval shown in bars. The figure shows panel local projection estimates of the response of the indicated price indices to the LS instrument $ShareNearCLL_{i,t} \times \%ChangeInCLL_t$ as estimated using equation (1). Control variables include $ShareNearCLL_{i,t}$ and its lag and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

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 then 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 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 l

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

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

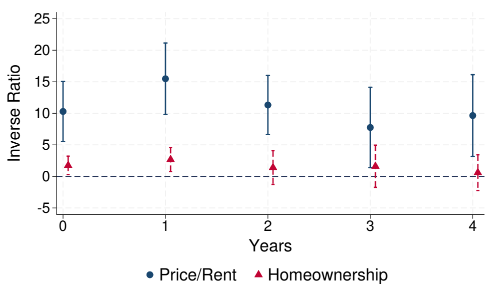

(a) Point Estimates (HVS)

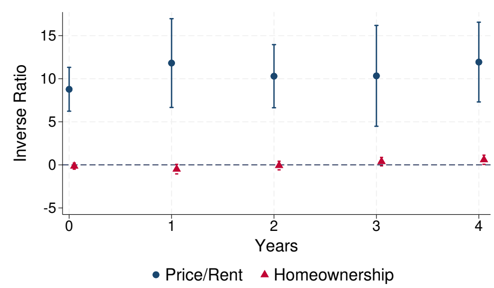

(b) Point Estimates (GG Microdata)

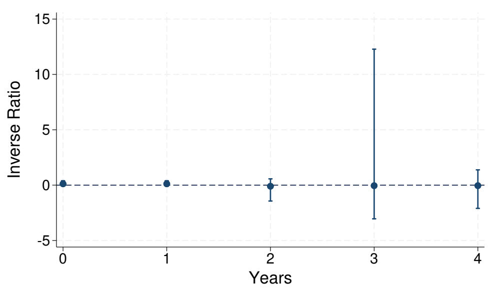

(c) Inverse Ratio (HVS)

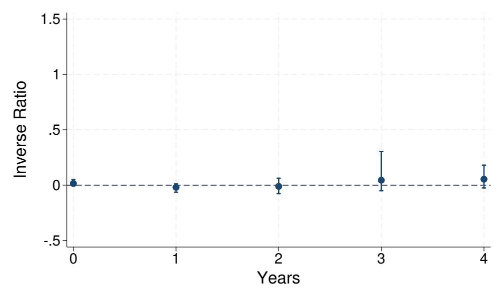

(d) Inverse Ratio (GG Microdata)

Notes: 95% confidence interval shown in bars. The figure shows panel local projection estimates of he 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

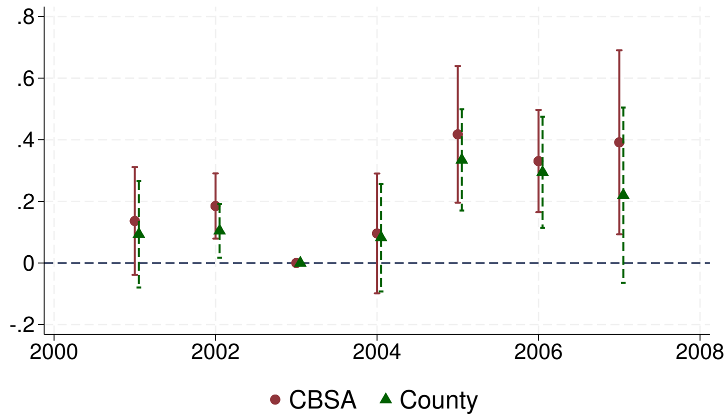

Notes: 95% confidence interval shown in bars. The figure shows estimates of the estimates of β_k for each indicated year and outcome variable (price or the GG microdata-based homeownership rate) estimated from equation (2), with the instrument being $Z_i = APL_{2004} \times OCC_{2003}$. The controls 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.

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

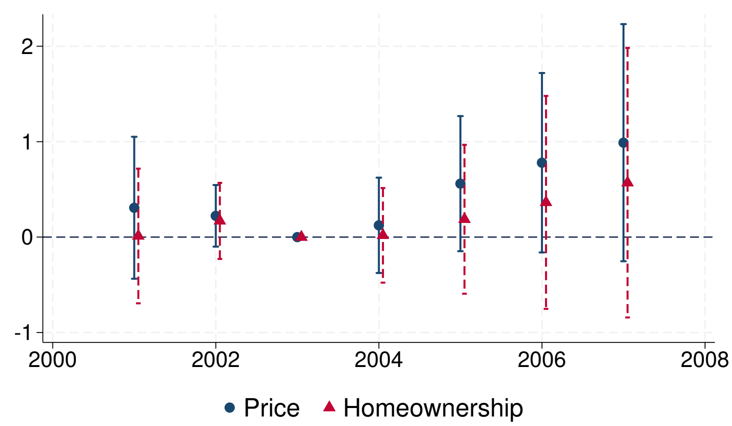

Notes: 95% confidence interval shown in bars. The figure shows estimates of the cumulative sum from 2003 of β_k for each indicated year and outcome variable (price or the GG microdata-based homeownership rate) estimated from equation (2), with the instrument being $Z_i = APL_{2004} \times OCC_{2003}$. 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

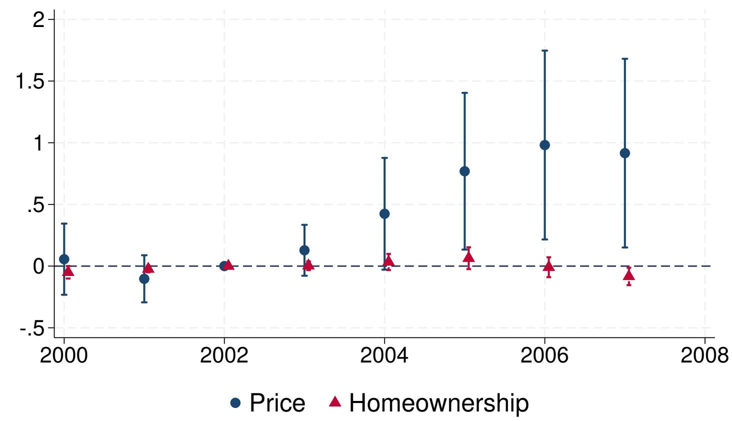

Notes: 95% confidence interval shown in bars. The figure shows shows estimates of the effect of a city's NCL share on the indicated outcome (price or the GG microdata-based homeownership rate) based on estimating equation (2) with the instrument being $Z_i = NCLShare_i^{2002}$ and 2002 being the base 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

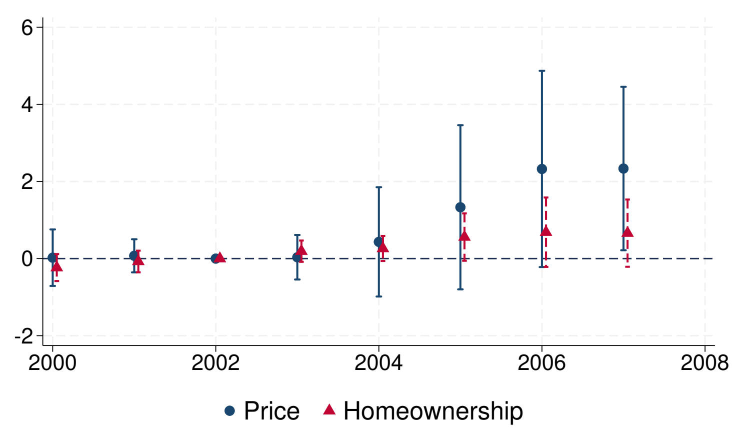

Notes: 95% confidence interval shown in bars. The figure shows shows estimates of the effect of a city's NCL share on the indicated outcome (price or the GG microdata-based homeownership rate) based on estimating equation (2) with the instrument being $Z_i = NCLShare_i^{2002}$ and 2002 being the base 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 our construction of the aggregate price-rent ratio, used in Figures 1, as well as for our main model experiments. Our measure of the national price-rent ratio comes from the BEA and the Financial Accounts of the United States 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 method of Chow and Lin (1971), which approximates a higher-frequency series using a lower-frequency version of that series and a higher-frequency related series. For the lower-frequency series (y) we use the log of the annual series for owner-occupied housing services. For our related or proxy series, we use the log of the product of total housing services and the national homeownership rate:

$$x_t = \log\Big(TotalServices_t \times HOR_t\Big)$$

For intuition, this series would be exactly equal to the target (log owner-occupied housing services) under the assumption that housing services per unit of housing is identical for owners and renters. We then create a quarterly version of the log owner-occupied housing services series (z_t) implied by the regression relationship

$$Y = BX + \varepsilon$$

where the matrix

$$B = \begin{bmatrix} 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & \cdots & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 & \cdots & 0 & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \cdots & 1 & 1 & 1 & 1 \end{bmatrix}$$

ensures that each annual element of y should be an average over four quarters of x.[7] We follow the formulas in Chow and Lin (1971) to compute the high-frequency proxy z_t, which we use as the quarterly value of log owner-occupied housing services. Taking the ratio of the market value of household real estate to the level of this value produces the desired price-rent ratio.

[7] In principle, the averaging relationship should hold in levels rather than logs. However, we choose to use logs to avoid issues with the changing scale of our variables over time.

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, also known as ZTRAX (Zillow, 2021). 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.

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 abaove.

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 occuped 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. We will denote this initial measure for county i and year t as $HOR_{i,t}^{GG,\text{raw}}$.

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 remote time effects from our homeownership measures using

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

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 $\overline{HOR}^{s}_{t}$ are computed as weighted averages of $HOR^{s}_{i,t}$ 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:

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

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:

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

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 $\widehat{Census}_{i,2010}$ 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 $\widetilde{HOR}^{GG,\text{raw}}$ 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 $\widetilde{HOR}_{i,d}^{GG,\text{raw}}$, where $d \in \{1980, 1990, 2000\}$ is a Decennial Census year. For example,

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

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

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

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

   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 $\widetilde{HOR}_{i,t}^{GG,\text{raw}}$, and subtract off the county-specific low frequency trend in our "GG, raw" data $\widehat{HOR}_{i,t}^{GG,\text{raw}}$. 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 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

Table C.1: Homeownership Rate Comparison

Statistic	ACS	HVS	GG
Standard deviation (levels)	0.70%	1.87%	0.34%
Standard deviation (1Y differences)	0.94%	2.18%	0.29%
Autocorrelation (levels)	0.089	0.331	0.738
Autocorrelation (1Y differences)	-0.395	-0.108	0.669
Standard deviation (deviation from ACS)	–	2.36%	0.94%
Correlation with ACS (levels)	1.000	0.058	0.216
Correlation with ACS (1Y differences)	1.000	0.027	0.092
Correlation with ACS (3Y differences)	1.000	0.036	0.245
Correlation with ACS (5Y differences)	1.000	0.081	0.339

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. Statistics are computed after removing CBSA and time fixed effects, as well as a linear time trend. The sample includes all CBSAs with at least ten years of data for the ACS, HVS, and GG series.

2005 to 2017.[8]

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.5. 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 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

[8]There 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.

Table C.2: Regression Results

	$ACS_{j,t}$	$HVS_{j,t}$	$GG_{j,t}$
$ACS_{j,t+1}$	0.051	0.002	0.392***
	(0.043)	(0.014)	(0.076)
N	885		
Adjusted R^2	0.036		

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.

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:

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

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, 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

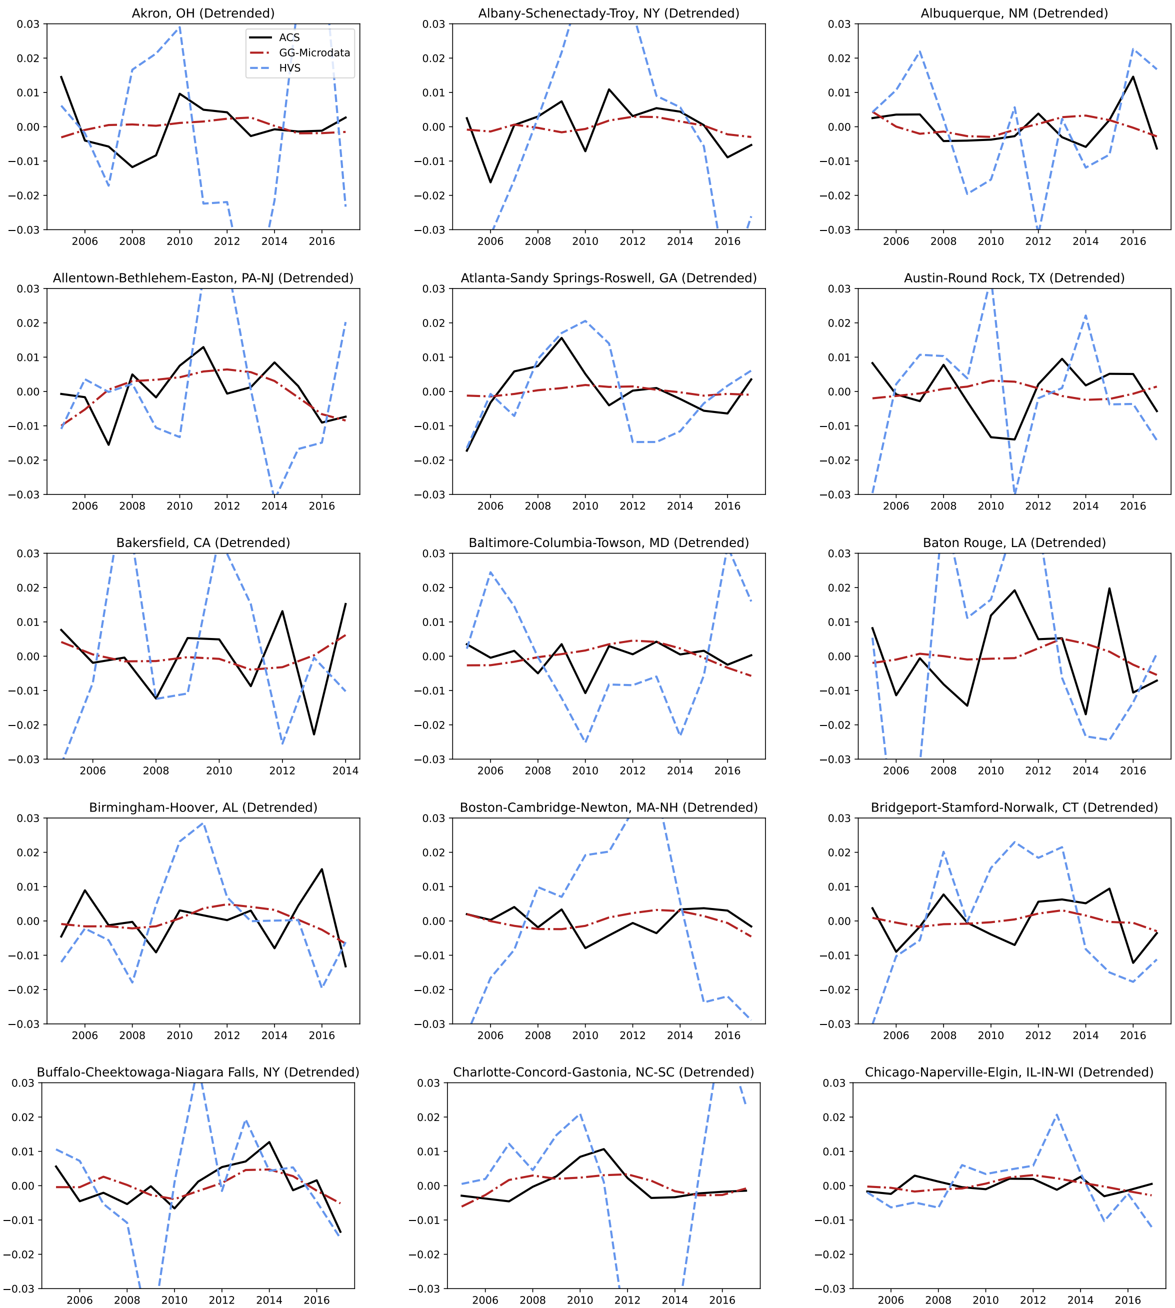

Notes: The series "ACS," "GG-Microdata," and "HVS," correspond to homeownership rates measured using the American Community Survey, our new microdata-based measure, and the Census Housing and Vacancy Survey. For each series, we have removed CBSA and year fixed effects, as well as a linear time trend.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

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

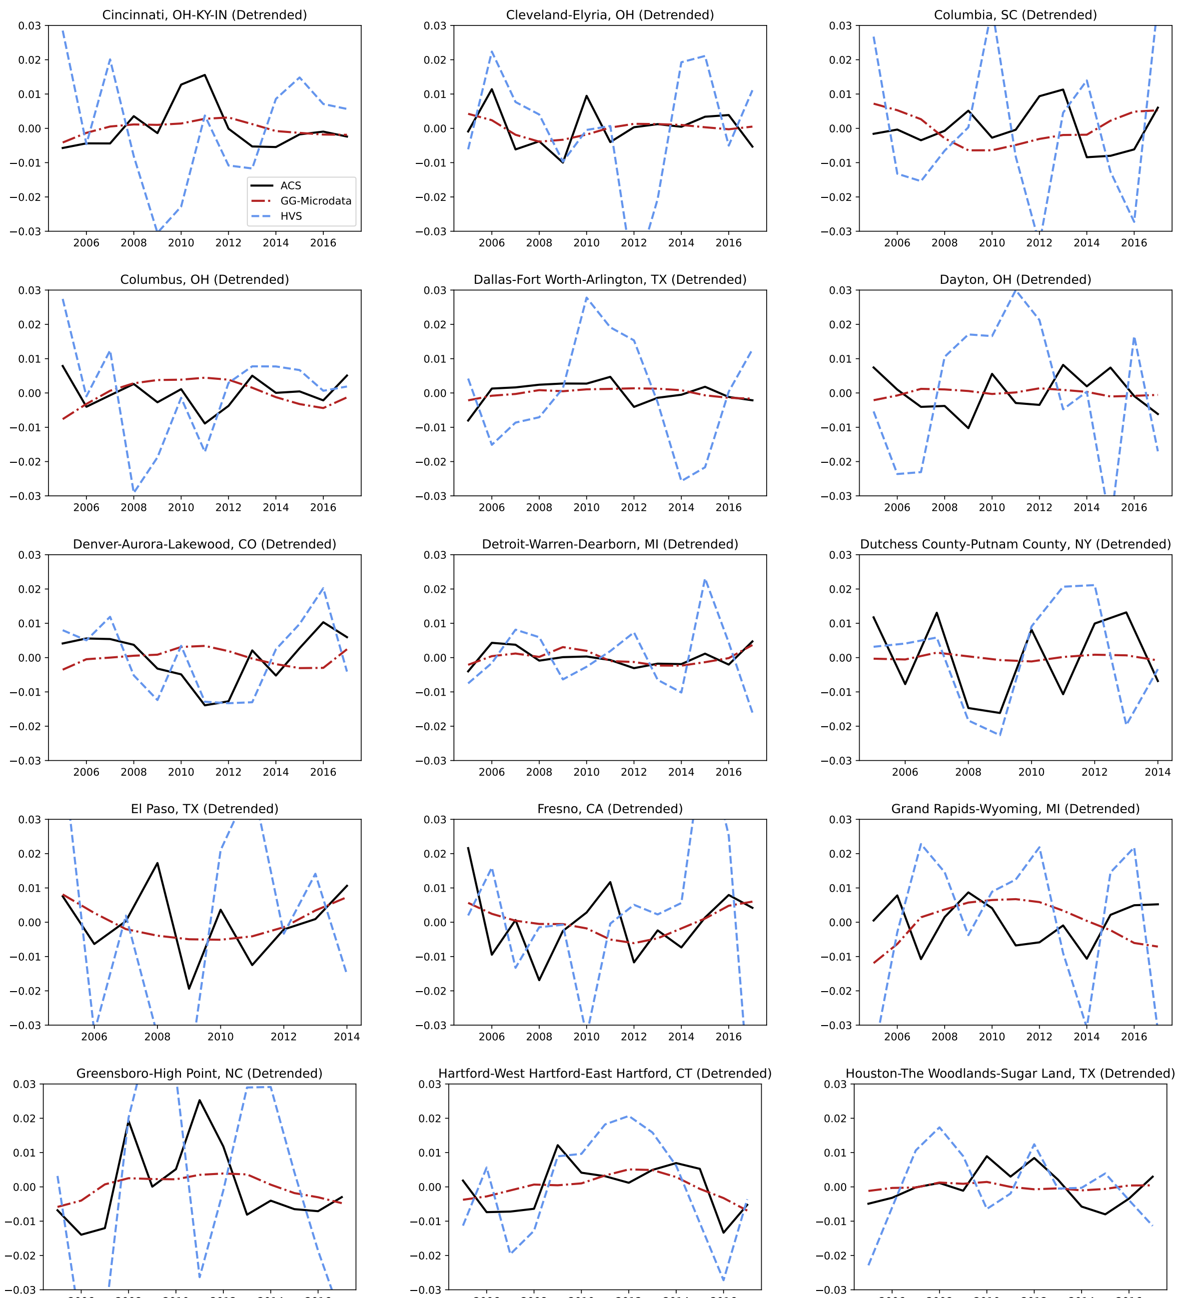

Notes: The series “ACS,” “GG-Microdata,” and “HVS,” correspond to homeownership rates measured using the American Community Survey, our new microdata-based measure, and the Census Housing and Vacancy Survey. For each series, we have removed CBSA and year fixed effects, as well as a linear time trend.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

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

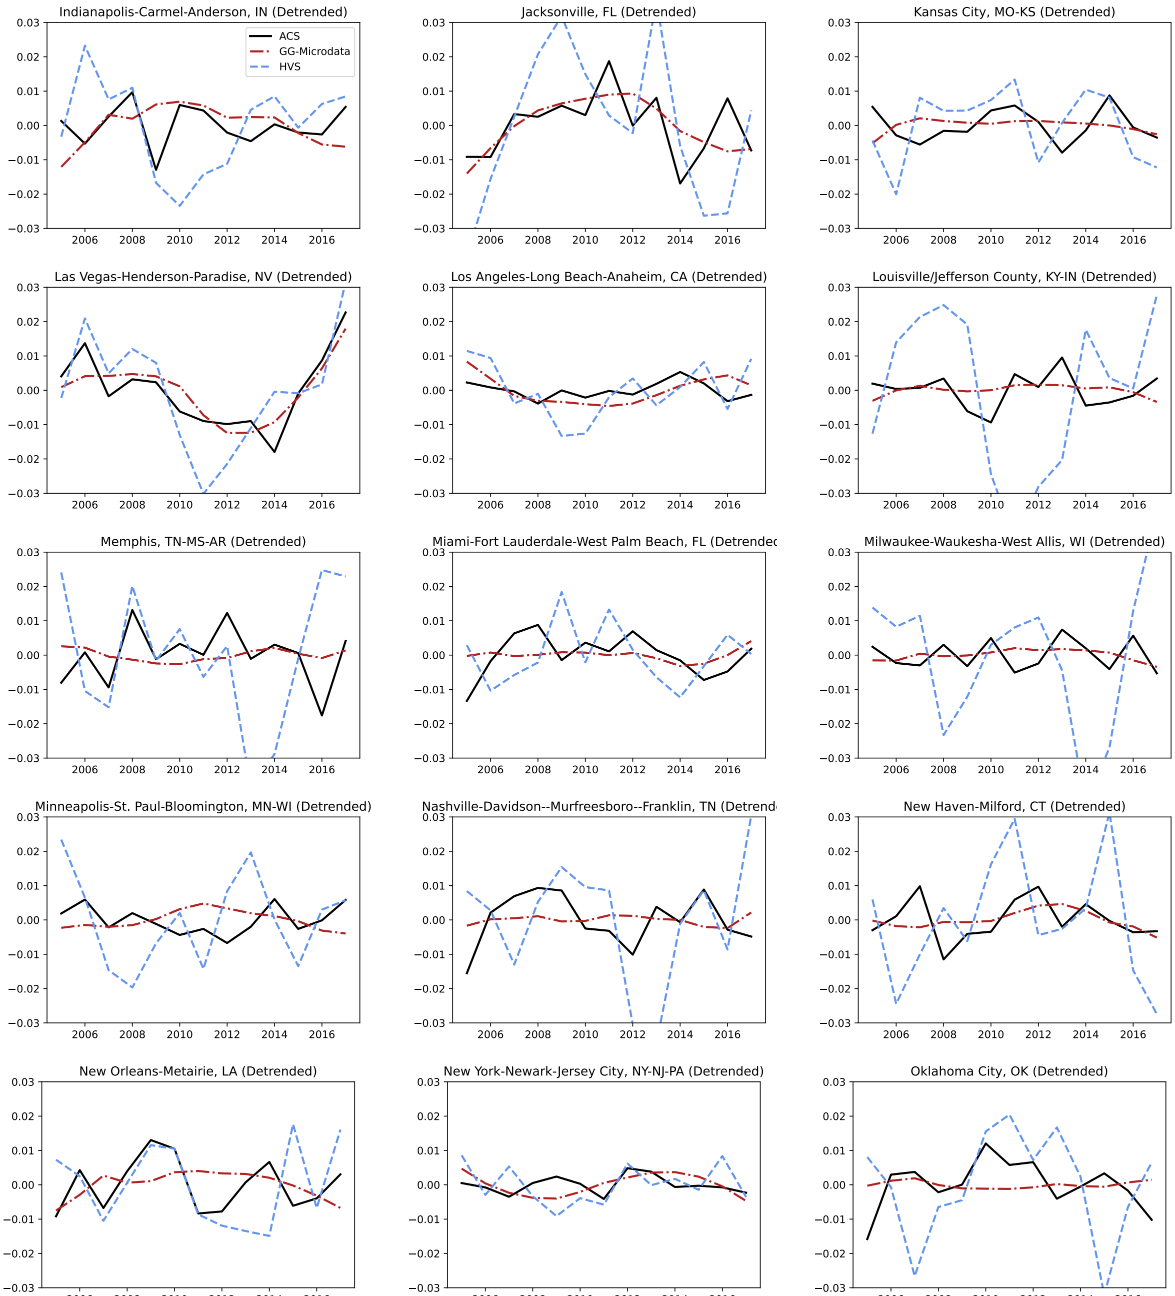

Notes: The series "ACS," "GG-Microdata," and "HVS," correspond to homeownership rates measured using the American Community Survey, our new microdata-based measure, and the Census Housing and Vacancy Survey. For each series, we have removed CBSA and year fixed effects, as well as a linear time trend.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

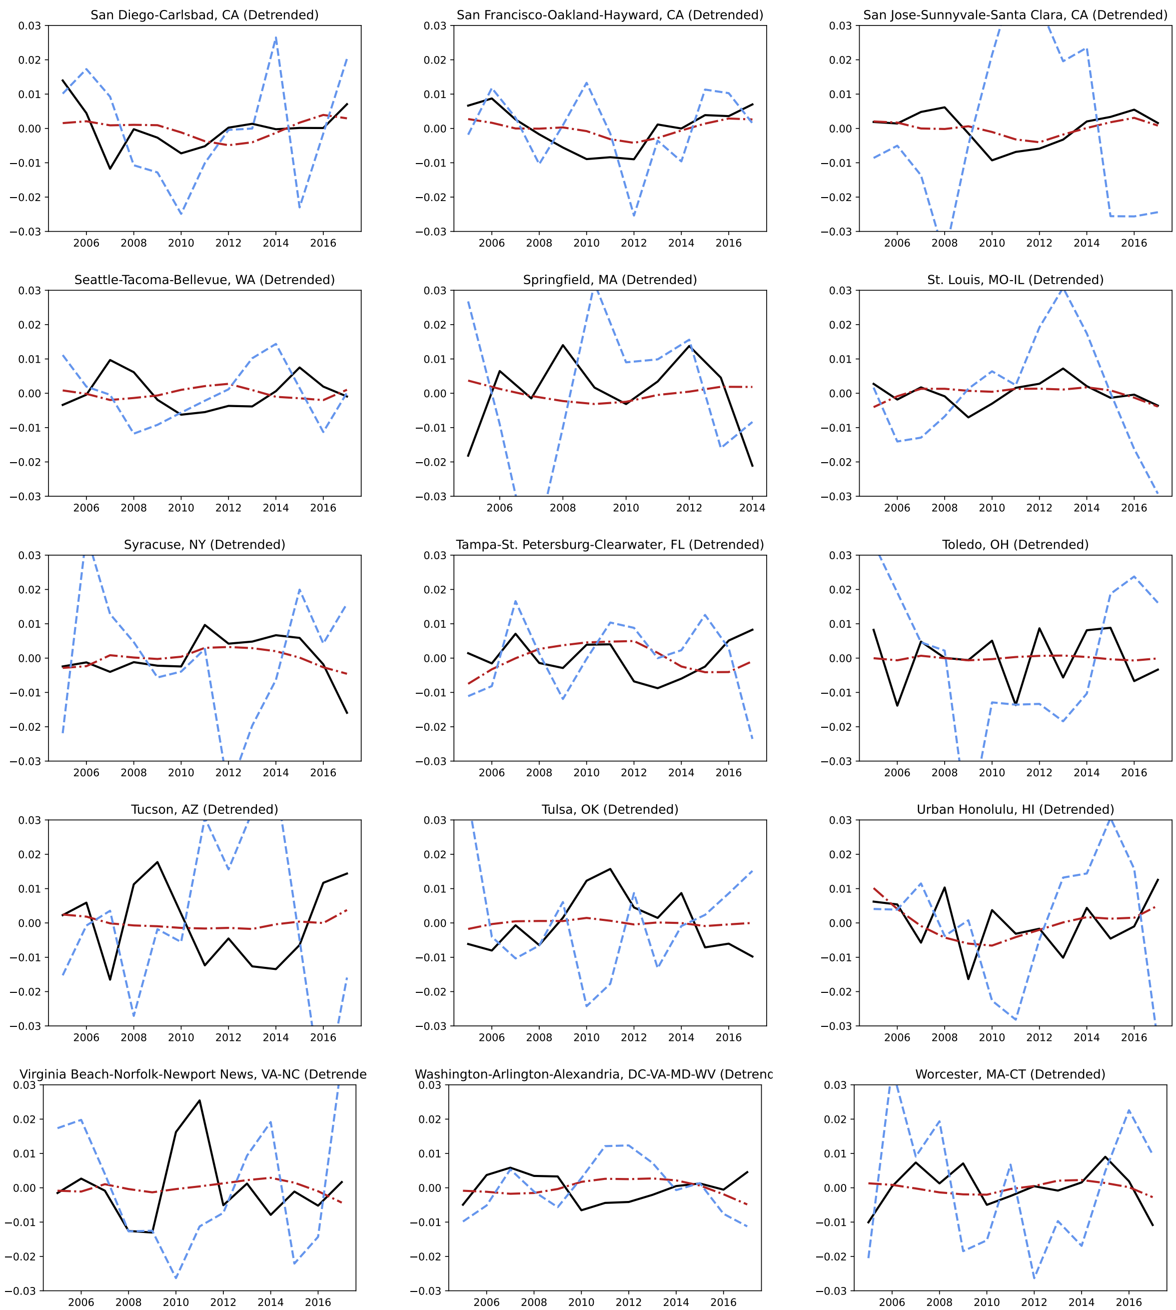

Notes: The series "ACS," "GG-Microdata," and "HVS," correspond to homeownership rates measured using the American Community Survey, our new microdata-based measure, and the Census Housing and Vacancy Survey. For each series, we have removed CBSA and year fixed effects, as well as a linear time trend.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

References

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.

Chow, Gregory C and An loh Lin, "Best Linear Unbiased Interpolation, Distribution, and Extrapolation of Time Series by Related Series," *The Review of Economics and Statistics*, 1971, pp. 372–375.

Greenwald, Daniel, "The Mortgage Credit Channel of Macroeconomic Transmission," 2018. MIT Sloan Research Paper No. 5184-16.

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.

Landvoigt, Tim, Monika Piazzesi, and Martin Schneider, "The Housing Market(s) of San Diego," *American Economic Review*, 2015, *105* (4), 1371–1407.

Loutskina, Elena and Philip E Strahan, "Financial Integration, Housing, and Economic Volatility," *Journal of Financial Economics*, 2015, *115* (1), 25–41.

Saiz, Albert, "The Geographic Determinants of Housing Supply," *The Quarterly Journal of Economics*, 2010, *125* (3), 1253–1296.

Zillow, "Zillow Transaction and Assessment Database (ZTRAX)," 2021. Proprietary Data Obtained From Vendor.