DO CREDIT CONDITIONS MOVE HOUSE PRICES?

APPENDICES

Daniel L. Greenwald
Adam Guren

WORKING PAPER 29391

A Model Appendix

A.1 Equilibrium Conditions

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

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

$$(h_{B,t}): \qquad q_t = (u^h_{B,t}/u^c_{B,t}) \tag{10}$$

$$(H^*_{B,t}): \qquad p_t = \frac{\mathbb{E}_t\left\{\Lambda_{B,t+1}\left[\bar{\omega}_{B,t} + q_{t+1} + \left(1 - \delta - (1-\rho_B)\mathcal{C}_{B,t+1}\right)p_{t+1}\right]\right\}}{1 - \mathcal{C}_{B,t}} \tag{11}$$

$$(M^*_{B,t}): \qquad 1 = \Omega^B_{M,t} + r^*_{j,t}\Omega^B_{X,t} + \mu_{B,t}, \tag{12}$$

where:

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

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

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

Equation (10) sets the rent equal to the marginal rate of substitution between housing services and consumption. Equation (11) specifies that at an interior solution, the price of housing must be equal to the present value of next period's service flow (the rent combined with the owner's utility bonus) plus the continuation value. Note that since p_t is the price of newly purchased housing that is about to be borrowed against, p_t includes the value of collateral services, which the borrower does not receive in periods when she does not refinance. Therefore, the continuation value is equal to the market value of housing net of maintenance costs, $(1-\delta)p_{t+1}$, minus the value of collateral services $\mu_{B,t+1}\theta_B$ in states of the world when the borrower does not refinance, which occurs with probability $(1-\rho_B)$. Equation (12) sets the marginal benefit of one unit of face value debt ($1 today) against the marginal cost (the continuation cost of the debt plus the shadow cost of tightening the borrowing constraint).

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

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

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

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

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

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

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, the landlord's budget constraint becomes:

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

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

while the landlord's laws of motion are:

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

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 optimality conditions are:

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

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

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

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

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

Calibration. To calibrate the model to match our empirical IRFs as in Section 6, we need to map our identified LS shock into the model. While this mapping into a shock to mortgage spreads is straightforward for the model with borrower credit only, as the vast majority of conforming mortgages are securitized by the GSEs, this is not the case for landlords. While single family rental properties, and multifamily rental properties with 2-4 units, can be financed by GSE loans, and are thus influenced by changes in the conforming loan limit, rental properties with more than four units are not eligible for this type of mortgage, and thus should not be directly affected by the LS shock. Since roughly 50% of rental units are in eligible 1-4 family buildings (Joint Center for Housing Studies of Harvard University (2020)), we choose the parsimonious recalibration $s_B = 1, s_L = 0.5$, so that half of landlord credit is eligible for the GSE subsidy.

A.3 Model: Additional Results

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

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

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

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

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

Figure A.1: Credit Relaxation by Borrower Heterogeneity

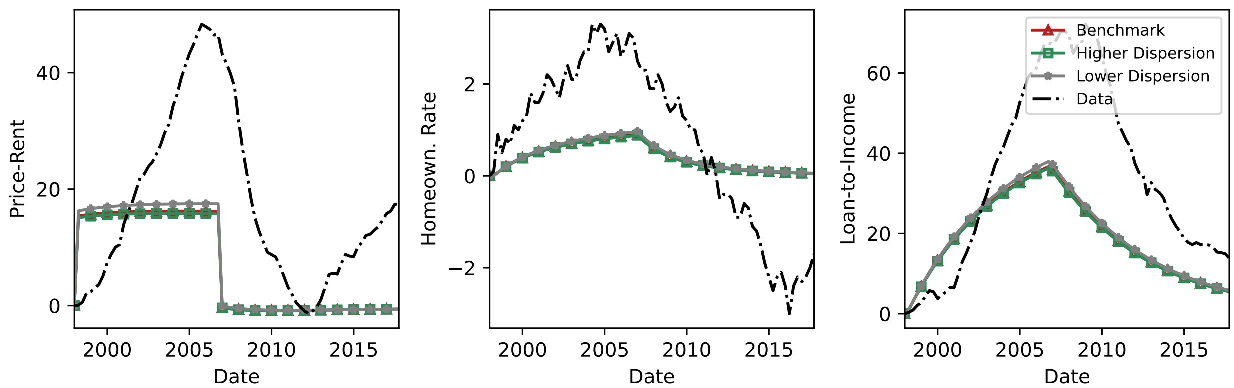

(a) Credit Relaxation by Borrower Heterogeneity

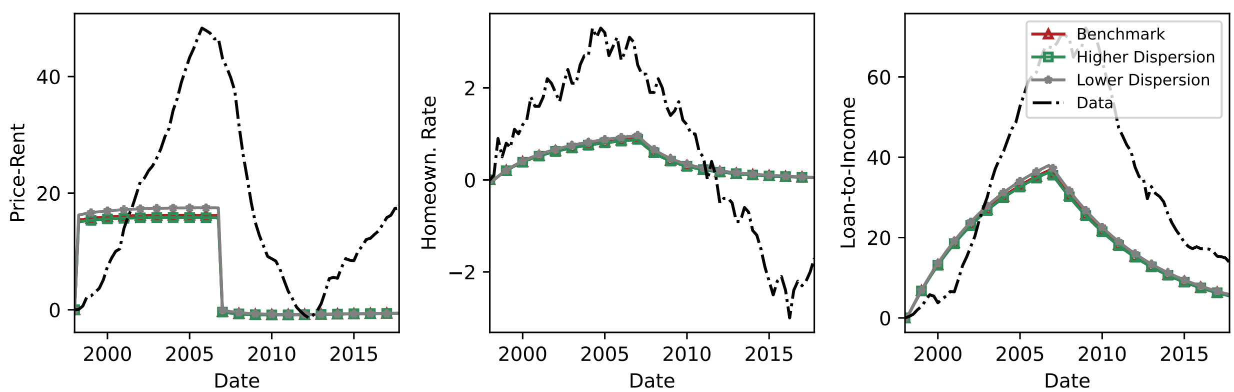

(b) Credit Relaxation by Borrower Heterogeneity

Notes: Panel (a) shows responses to our Credit Relaxation experiment from Figure 7a varying the level of borrower dispersion (σ_B). Panel (b) repeats this exercise for our Credit Relaxation + Rates experiment from Figure 7b. The "Higher Dispersion" series sets $\sigma_{\omega,B}$ so that half as many renters ($0.64\%/2 = 0.32\%$) switch to ownership under the First Time Homeownership Subsidy, while the "Lower Dispersion" series sets $\sigma_{\omega,B}$ so that twice as many renters ($2 \times 0.64\% = 1.28\%$ switch. The loan-to-income ratio is the ratio of household debt (FRED code: HMLBSHNO) to household gross income (FRED code: PI) in the Flow of Funds. For other data definitions see the notes for Figure 1.

B Empirical Appendix

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

B.1 Data Construction

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

B.1.1 LS Instrument Data Set

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

Our data sources are:

- House Prices: Our primary data source 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 from 1976 to 2018. We use FHFA house price indices and Freddie Mac house price indices at the CBSA level as a supplementary data source.

- Rents: Our rent series is the CBRE Economic Advisers Torto-Wheaton index. In particular, we use their nominal rent index. This is available for 66 geographic areas that we map to CBSAs from 1980 to 2017. We are able to map to a quarterly panel for 53 CBSAs beginning in 1989 and 62 CBSAs beginning in 1994.

- Homeownership rates:
  - Our main homeownership measure is from the Census Housing Vacancy Survey. The Census produces homeownership rates at the CBSA level, however the CBSA definitions change over time. They use 1980 MSA definitions from 1986-1994, 1990 MSA definitions from 1995-2004, 2000 CBSA definitions from 2005-2014, and 2010 CBSA definitions from 2015-2017. We use a crosswalk to link these longitudinally. To deal with changing definitions, we use data on homeownership rates aggregated from the county level to each MSA/CBSA definition. If the difference between the homeowner rates using the two different CBSA definitions is more than 5% in either of the two closest censuses, we drop all log changes that go across the definition break point from our analysis. For instance, the 1990 MSA definitions include far fewer suburbs in the New York Metropolitan Area than the 2000 CBSA definition. As a result, in the Census data the homeownership rate is 37% in 2004 and 55% in 2005. Using the Census data, we see that using both the 2000 and 2010 census data (the two closest censuses to the 2004-2005 switch), the difference between the homeownership rates based on the two definitions is over 45%, so we drop any log changes in the homeownership rate that cross over the 2004-2005 redefinition. For robustness, we add analysis at the state level where geographic definitions do not change.
  - Our secondary homeownership measure is from the ACS one-year estimates at the CBSA level, which is downloaded from the Census.
- Credit Data: Our credit data come from the Home Mortgage Disclosure Act microdata, which we collapse to the CBSA level. Our main measure is the dollar volume of loans originated, but in the appendix we use robustness to the number of loans

and the loan to income ratio. We also use the HMDA data together with the history of national conforming loan limits from the FHFA to create our instrument, which is the fraction of originations within 5% of the CLL.

- Housing Units: Our housing units data come from the Census.
- Housing supply elasticity: We use data from (Saiz) which we crosswalk from his MSA definitions to our CBSAs using principal cities.
- Employment and industry shares: We use the annual series of county-level employment from the QCEW and aggregate to the CBSA level to create a measure of log employment and employment shares for each NAICS two-digit industry. The QCEW suppresses observations where employment in a county-year-industry is small. To handle cases where a county barely slips below the suppression threshold for one year, we linearly interpolate employment when we have a few missing years. For other cases, employment is small enough for a missing year that ignoring the issue does not matter once we aggregate to the CBSA level.

One data source that merits additional discussion is the CBRE Torto-Wheaton rent index. As mentioned in the main text, it 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.

B.1.2 DK Instrument Data Set

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

B.1.3 MS Instrument Data Set

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

B.2 LS Instrument Details and Robustness

The Loutskina-Strahan instruments are the interaction of the change in the national conforming loan limit and the share of HMDA mortgage originations within 5% of the conforming loan limit in the prior year, as well as the interaction of this first instrument with the Saiz housing supply elasticity. We use the CLL for single-unit mortgages provided by FHFA. As mentioned in a footnote in the main text, starting in 2008 Congress allowed the

CLL to rise by more in high-cost cities if their local house price index grew sufficiently quickly. This would violate an instrumental variable's exclusion restriction because the change in the CLL would be mechanically correlated with lagged local outcomes. Consequently, in constructing the instrument we use the change in the national CLL regardless of the change in the local CLL in high-cost areas.

B.2.1 First Stage

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

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

B.2.2 Additional Results and Robustness

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

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

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

Table B.1: Loutskina-Strahan First Stage

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

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

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

Figures B.3 and B.4 show results when we use the number of loans and the loan-to-income ratio as the credit measure rather than the dollar value of loans, following Favara

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

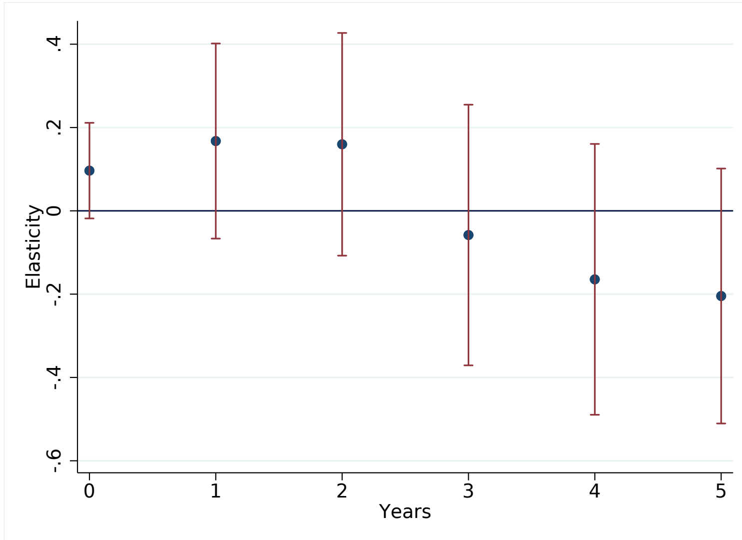

Notes: 95 % confidence interval shown in red bars. The figure shows panel local projection instrumental variables estimates of the response of log rents to dollar credit volume. The second stage is as indicated in equation (1) and the first stage is as indicated in equation (2). The two instruments are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

and Imbs (2015). Things look broadly similar in shape but not in magnitude. This makes sense: the reduced form between the instrument and the outcome remains the same, but the first stage by which it is scaled changes as we change the X variable.

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

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

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

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

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

Finally, Figures B.8 and B.9 show results using slightly different data than our main analysis. Figure B.8 uses the FHFA purchase-only house price index rather than the CoreLogic house price index. One can see a similarly shaped IRF, although the magnitudes are larger with the FHFA index. As a result, the inverse supply slope is similar, although slightly smaller. Figure B.9 uses ACS homeownership rate data rather than Census. The

ACS has the advantage of having consistent geographies but it only starts in 2005, so we lose a lot of the variation in the instruments in the boom (there is very little variation in the bust since the conforming loan limit only adjusts upwards and was thus flat throughout the bust). The standard errors are thus quite wide, especially at the 2-year horizon, but the point estimates are generally of a similar magnitude. The inverse slope impulse response is similar although slightly smaller in magnitude until year 3, after which the estimates are imprecise. Overall, our results are robust to the data series used.

B.3 DK Instrument Details and Robustness

As described in the main body, we use a county-level data set generously provided by Di Maggio and Kermani which we collapse to the CBSA level. We are able to run the same regression at the CBSA level that they run at the county level with the exception of one proprietary control variable: the share of loans that are subprime. We also use log changes rather than percent changes. We run the same regression as DK, then transform the impulse response from log changes to log levels by cumulating the coefficients from the log changes regression.

B.3.1 Replicating DK's Exact Specification

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

B.3.2 Robustness

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

B.4 MS Instrument Details and Robustness

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

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

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

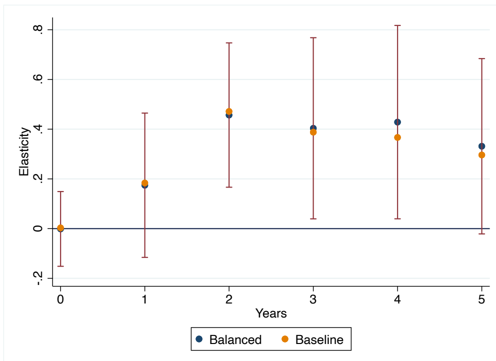

(a) Price/Rent

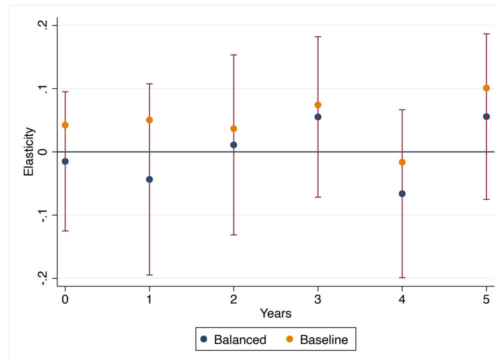

(b) Homeownership Rate

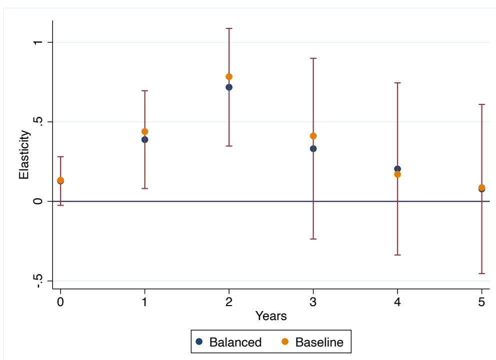

(c) House Price

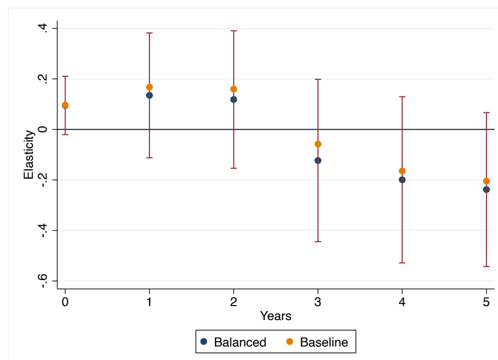

(d) Rents

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

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

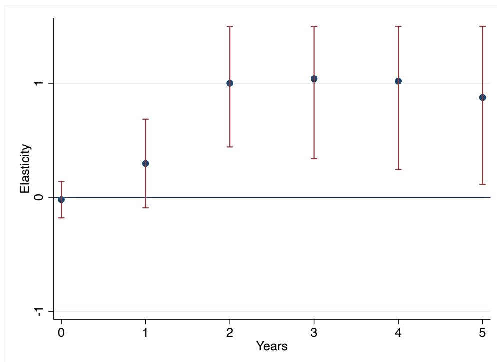

(a) Price/Rent

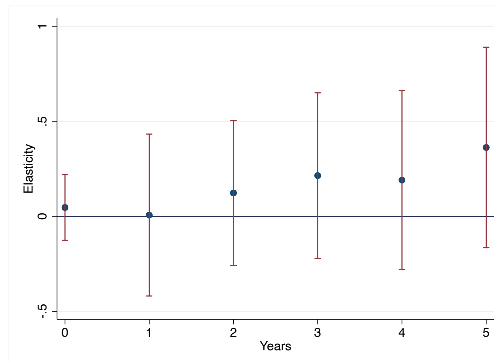

(b) Homeownership Rate

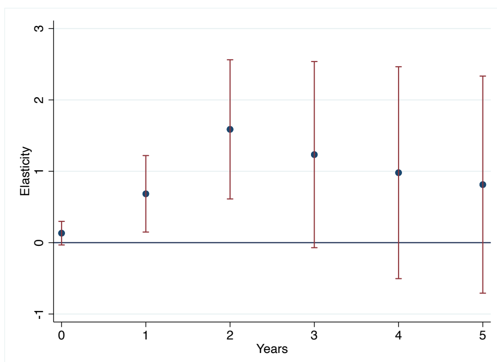

(c) House Price

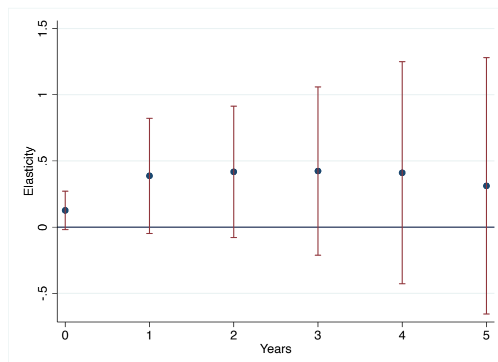

t

(d) Rents

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

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

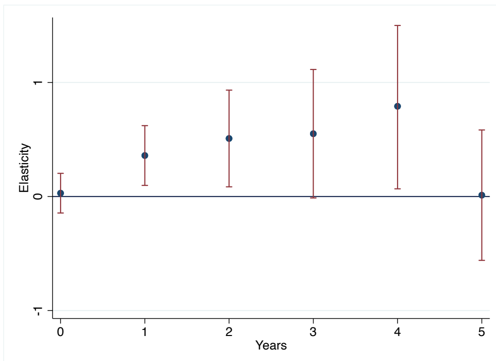

(a) Price/Rent

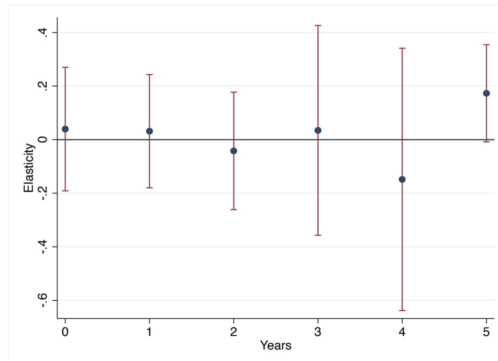

(b) Homeownership Rate

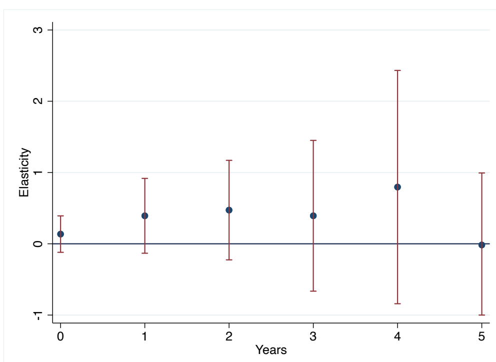

(c) House Price

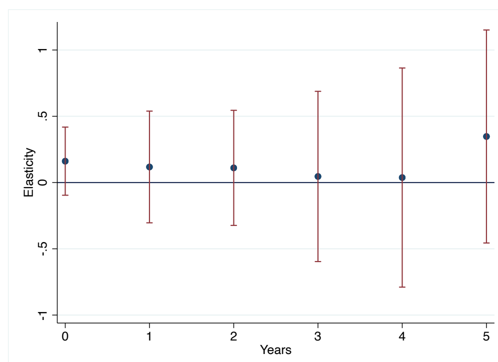

(d) Rents

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

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

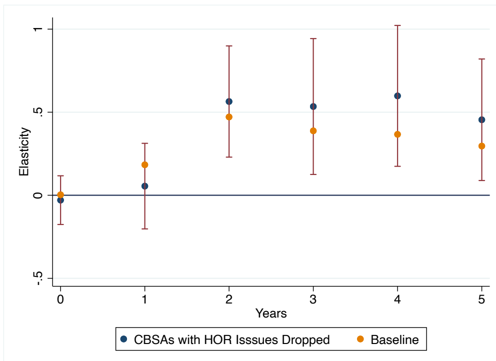

(a) Price/Rent

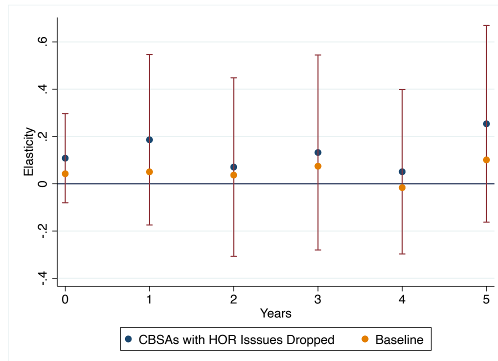

(b) Homeownership Rate

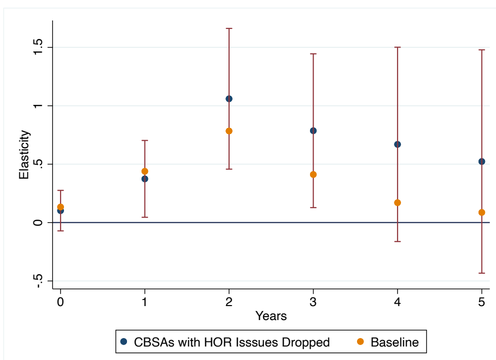

(c) House Price

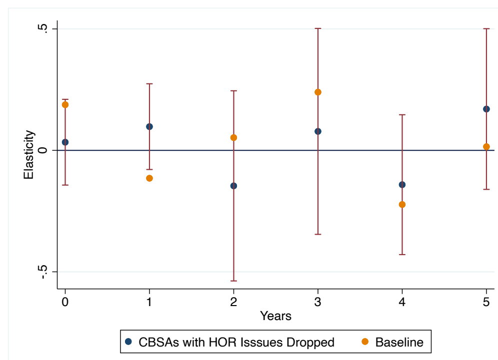

(d) Inverse Supply Slope

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

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

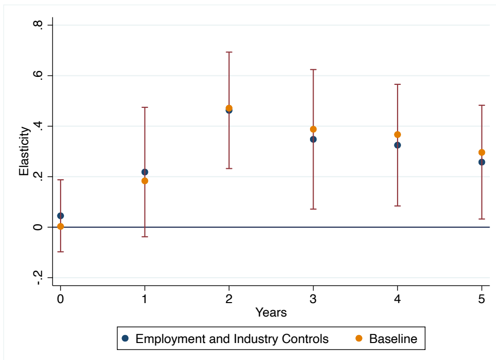

(a) Price/Rent

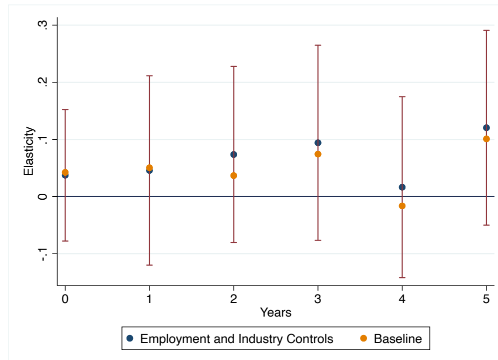

(b) Homeownership Rate

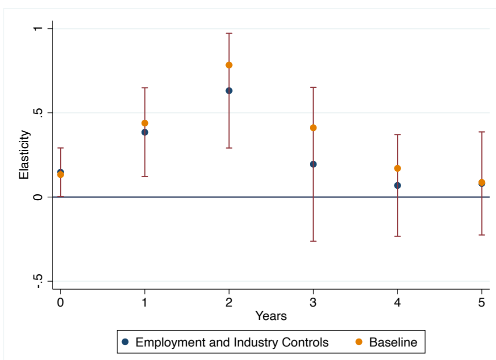

(c) House Price

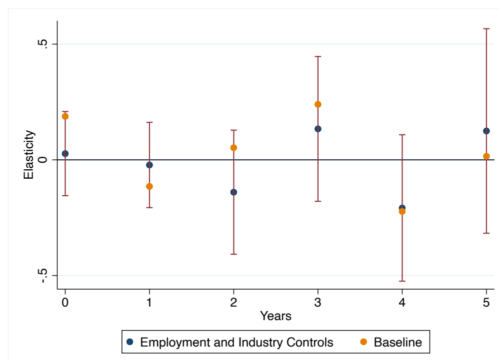

(d) Inverse Supply Slope

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

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

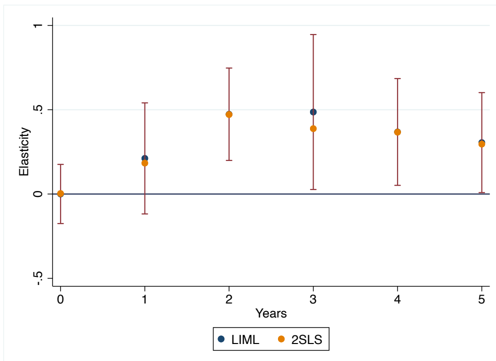

(a) Price/Rent

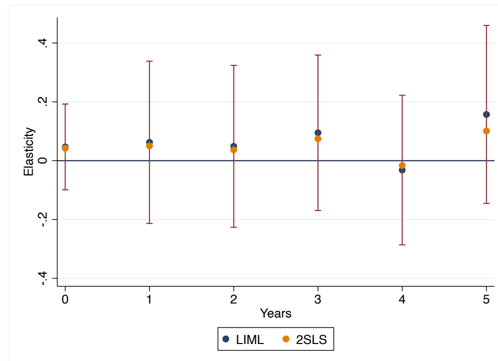

(b) Homeownership Rate

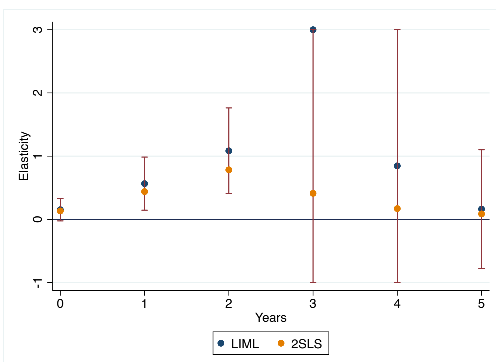

(c) House Price

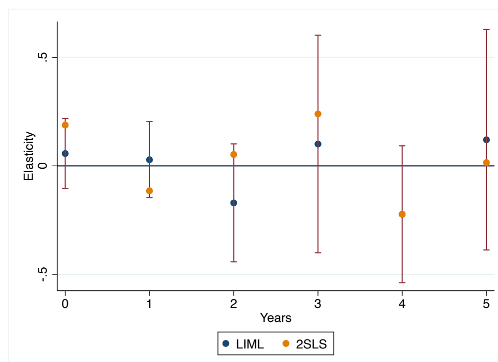

(d) Inverse Supply Slope

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

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

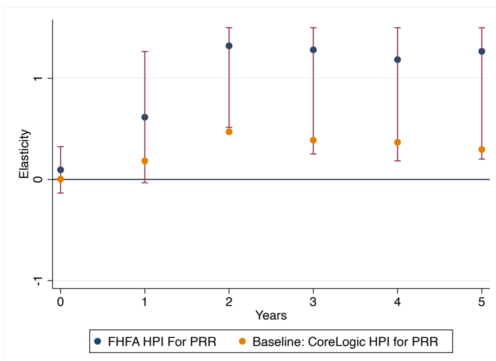

(a) Price/Rent

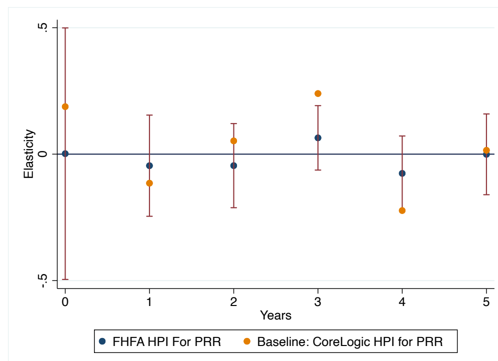

(b) Inverse Supply Slope

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

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

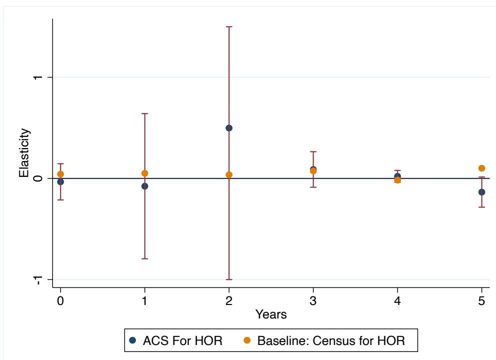

(a) Homeownership Rate

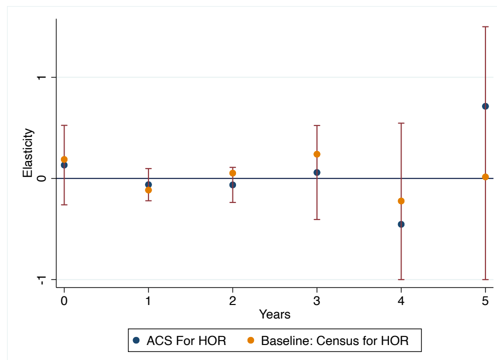

(b) Inverse Supply Slope

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

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

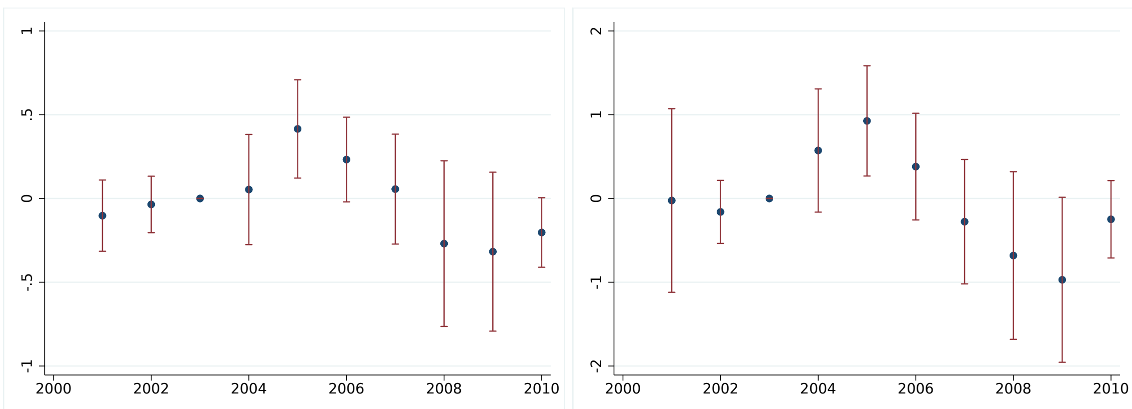

(a) House Price Growth (b) House Price Growth (HOR Sample)

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

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

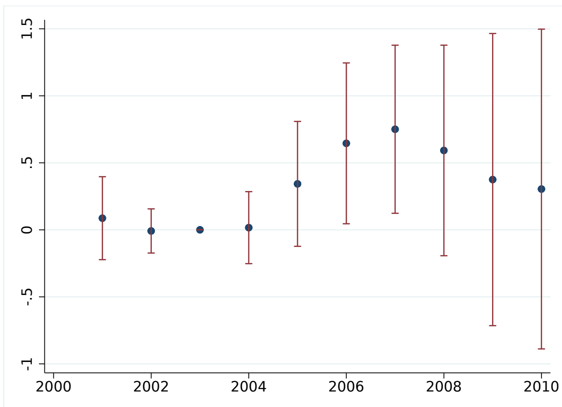

(a) House Prices

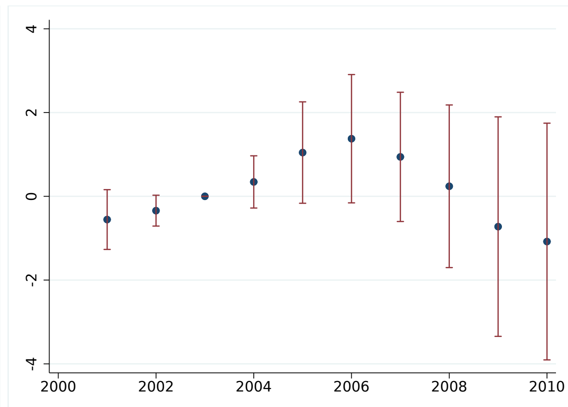

(b) House Prices (Homeownership Sample)

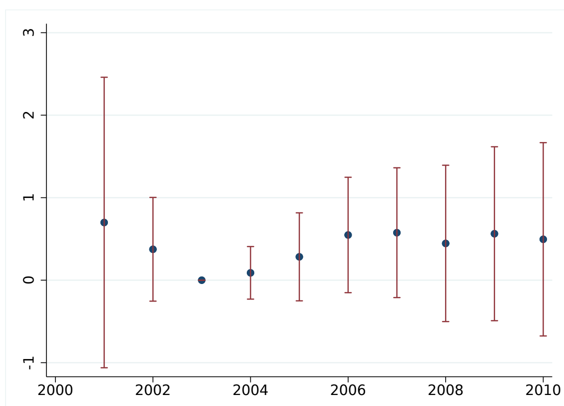

(c) Homeownership Rate

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

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

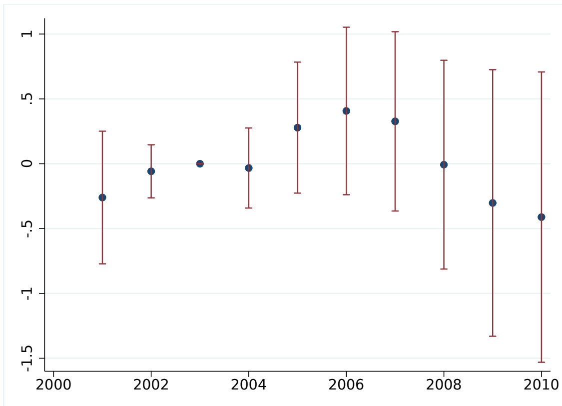

(a) House Prices

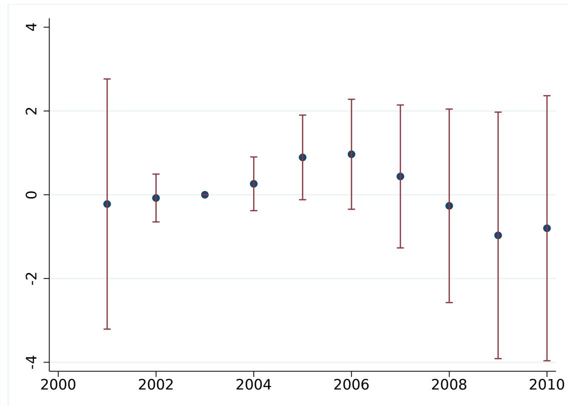

(b) House Prices (Homeownership Sample)

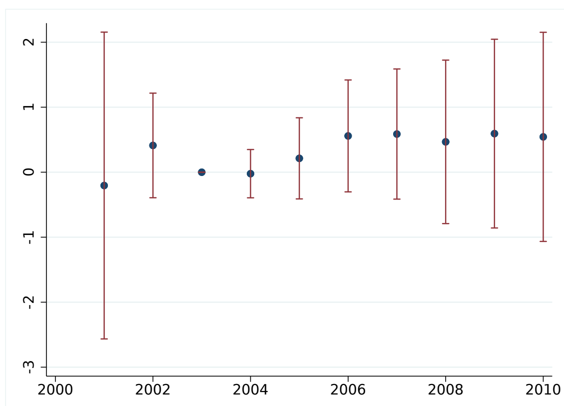

(c) Homeownership Rate

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

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

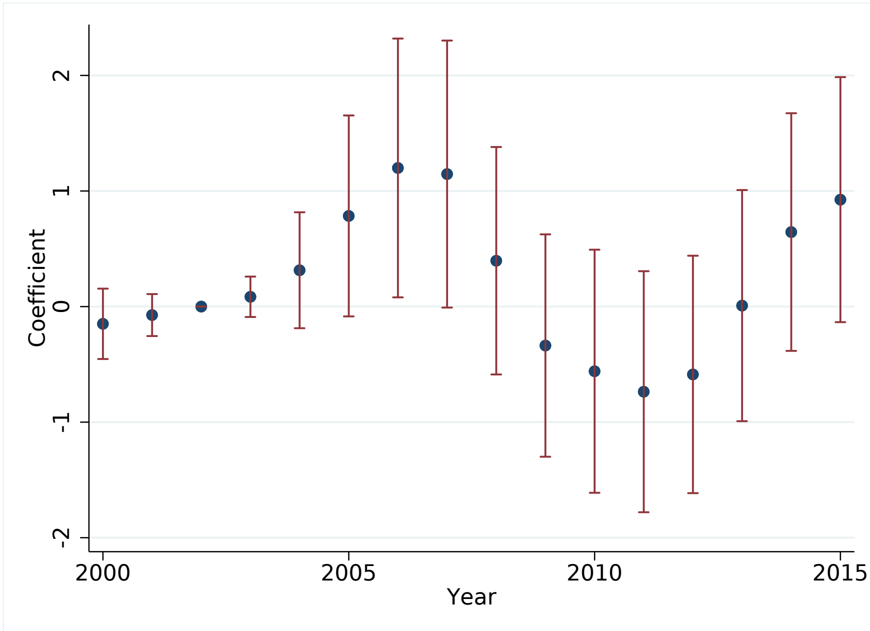

Notes: 95 % confidence interval shown in red bars. Figure shows estimates of the effect of a city's NCL share on each outcome based on estimating equation 3 with the instrument being $Z_i = NCLShare_i^{2002}$ and 2003 being the base year. All standard errors are clustered by CBSA as in the MS paper.

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

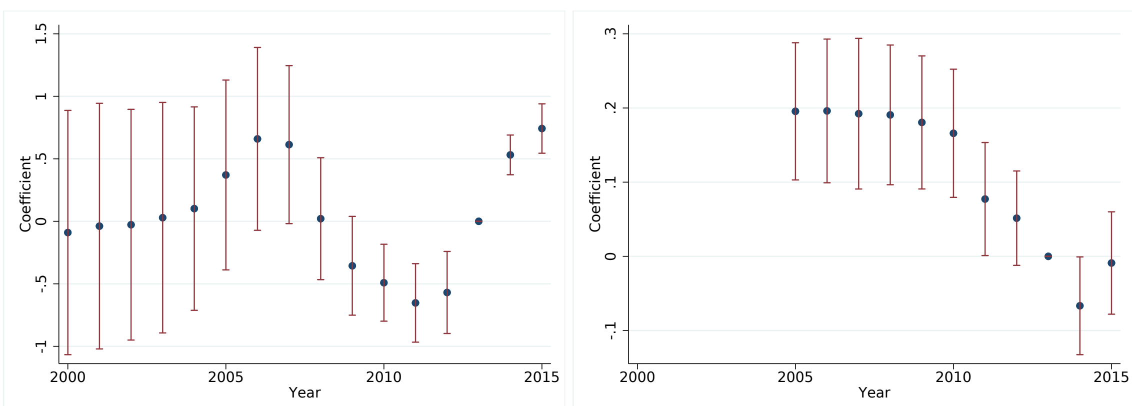

(a) House Prices

(b) Homeownership Rate

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

B.5 National Price-Rent Ratio Construction

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

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

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

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

where the hat denotes the interpolated value of this ratio.