University of Essex

Department of Economics

Discussion Paper Series

Do Households Use Homeownership To Insure Themselves? Evidence Across U.S. Cities

Michael Amior and Jonathan Halket

Do Households Use Homeownership To Insure Themselves? Evidence Across U.S. Cities

Michael Amior[*]and Jonathan Halket[†]

July 28, 2012

Abstract

Are households more likely to be homeowners when "housing risk" is higher? We show that homeownership rates and loan-to-value (LTV) ratios at the city level are strongly negatively correlated with local house price volatility. However, causal inference is confounded by house price levels, which are systematically correlated with housing risk in an intuitive way: in cities where the land value is larger relative to the local cost of structures, house prices are higher and more volatile. We disentangle the contributions of high price levels from high volatilities by building a life-cycle model of homeownership choices. The model is able to explain much of the cross-city dispersion in homeownership and LTV. We find that higher price levels explain the lower homeownership, while higher risk explains the lower LTV in high land value cities. The relationship between LTV and risk highlights the importance of including other means of incomplete insurance in models of homeownership. Finally, we use the model to show why regression-based inferences about the effect of risk on homeownership are biased.

Keywords: Homeownership, Housing Risk, Land Share, Loan-to-value, Life-Cycle

JEL Classification: D91, E21, R21, R31.

1   Introduction

Are households more likely to own their home when "housing risk" is higher? There is a large literature on how homeowners use home equity to smooth the transmission of earnings shocks into

[*]University College London, *m.amior@ucl.ac.uk*

[†]University of Essex and Institute for Fiscal Studies, corresponding author: *Economics Department, University of Essex, Wivenhoe Park, Colchester, CO4 3SQ, United Kingdom; jonathan@halket.com, http://halket.com/ Phone: +44 786-650-3704*

consumption[1]. In this paper, we explore how the decision to become a homeowner is influenced by exposure to housing market risk and motives for insurance.

There is growing evidence that households may bring forward their home purchase as a hedge against future house price fluctuations. But, an earlier purchase would often necessitate a larger mortgage and increased risk to consumption. Using a life-cycle model, we argue that, in response to differences in housing risk, otherwise similar households are more likely to differ in their liquid savings than the timing of their ownership decision. Empirically, these savings differences manifest themselves in observed variation in loan-to-value ratios (LTV). In other words, in response to higher price risk, households do not bring forward ownership decisions (nor do households that never own choose to become homeowners) - but rather, conditional on owning, they reduce their LTV.

We present supporting evidence from variation across US MSAs: cities with higher house price volatilities have lower average LTVs and also lower homeownership rates[2]. Based on estimates from our model, we argue these lower ownership rates are a consequence of the higher price levels in these risky cities. We also show that the high cross-sectional correlation between price levels and volatilities combined with the measurement error inherent in using historical volatility as a proxy for expected risk can lead to biases in reduced-form inferences of the effect of volatility on homeownership.

Theoretically, when markets are incomplete there are several reasons why homeownership may be a peculiar and attractive form of insurance against certain risks in the housing market[3]. In Ben-Shahar (1998), Nordvik (2001) and Sinai and Souleles (2005), households may use home-ownership to insure themselves against the risk of changes to the local rental price (or user-cost) of housing. However, Ortalo-Magne and Rady (2002) suggests that if a household's expected future earnings are more strongly correlated with local house prices, then it already has partial insurance through their labor earnings.[4] Ortalo-Magne and Rady (2006) and Banks et al. (2010) propose and

---

[1]Hurst and Stafford (2004) and Hryshko et al. (2010) look at how households use mortgage refinancing decisions and home equity, respectively, to smooth unemployment shocks and earnings shocks. Leth-Petersen (2010) finds that household expenditure increases moderately after credit constraints are relaxed. Paciorek and Sinai (2012) finds that the cross-sectional variance of housing consumption is lower for homeowners that have moved between cities whose house prices are strongly correlated.

[2]Throughout we refer to Metropolitan Statistical Areas as "cities" and "LTV" always refers to the loan-to-value at origination (that is, at the time of purchase). We will sometimes refer to the time-series standard deviation of the annual changes to log house prices within a city as its "volatility".

[3]Formal, direct means to insure against changes in house prices are limited (Caplin et al., 1997), and the correlation between house prices and other financial assets is small (Flavin and Yamashita, 2002).

[4]Davidoff (2006) finds that households purchase less housing when they work in an industry whose workers' income are relatively more correlated with local house prices. However, he finds very small effects of the same on the probability of homeownership.

find supporting evidence for a housing-ladder theory in which households who plan on eventually owning a large house (in part because larger houses may not be available on the rental market) are more likely to own a smaller home (rather than rent) first if they live in a risky area. This is to partially insure themselves against increases in the price of a good in their future consumption bundle (the larger house).

If financial constraints prevent some households from insuring themselves through owning, there may be important welfare improvements from policies designed to make ownership "accessible." However, measuring the size or even the overall sign of the insurance motive on homeownership is challenging, in part because it is difficult to isolate differences in households' exposure to housing risks that are independent from other factors that affect their homeownership decisions.

This study proceeds in three steps. First, we show that a substantial amount of the cross-city variation in housing risk is systematically related to the cross-city variation in the level of prices. We also show that households behave differently in the high-risk, high-price cities: they are less likely to become homeowners and more likely to make a large down payment (in percentage-of-house-value terms) when they do buy. So, secondly, we use a quantitative life-cycle model with homeownership to disentangle the effects of higher risk from higher price levels on the life-cycle timing of homeownership and mortgage decisions. Lastly, we show that typical regression-based inference procedures are biased measures of the effects of risk on homeownership because they use the ex-post, realized volatility of prices as their measure of ex-ante housing risk.

We document that both homeownership rates and average household loan-to-value (LTV) ratios at origination at the city level, controlling for household characteristics, are strongly negatively correlated with house price volatility. We measure local price volatility as the standard deviation of annual house price growth rates in the city over the previous five years; this is our "ex-ante" measure of housing risk[5]. In 2000, ownership rates were 60-80 percent for cities with the lowest price risk, but only about 30-50 percent for cities with the highest. And, for mortgages originating in 2000, mean LTV was more than five percentage points higher in the least risky cities compared with the riskiest.

However, causal inference is confounded by house price levels, which are systematically correlated with housing risk in an intuitive way: in cities where the land value is larger relative to the local cost of structures, house prices are higher and more volatile. When we look at the variation in homeownership rates and LTV by land share (the ratio of local land values to total housing costs), we see the same strong negative relationship. This is true even after instrumenting for possible

---

[5]"Ex-ante" in the sense that it is our measure of risk before considering the availability of assets that may provide for some degree of partial insurance.

endogeneity, using a measure of physical local land scarcity constructed by by Saiz (2010).

We focus on the cross-sectional dimension, rather than the time-series, for several reasons. For one, in the data, the amount of heterogeneity both in household and price behavior is much larger across-cities than within cities over time. For another, the land scarcity instrument offers a way to measure the effect of higher price levels and risk (jointly) on household behavior in the cross-section. However we cannot use it to control for endogeneity within a city over time. In this sense, our work is complementary to much of the previous literature (e.g. Sinai and Souleles (2005)), which exploits within-city variation in volatility (controlling for prices) to measure the insurance motive.

In order to measure the effect of higher volatility on homeownership, we disentangle its impact from that of higher prices by building a life-cycle model of homeownership choice. We account separately for innovations to house prices which are correlated with city wages and those which are not. The model has a flexible housing ladder where medium-sized housing can either be rented or owned, which enables the model to match the relative consumption of owner-occupied to rental housing in the average city according to land scarcity. Importantly, households have another means of imperfectly insuring themselves in addition to homeownership: a risk-free bond.

In our setup, households have several potential reasons why they might use homeownership to insure against housing risk. They may use homeownership to insure themselves against the risk of changes to the rental price of housing, though their labor earnings will provide some partial insurance. Also, the housing ladder assumption forces households that wish to live in large houses to own them so that the model nests the theories of Banks et al. (2010) and Ortalo-Magne and Rady (2006). Otherwise, the basic elements of our life-cycle model of homeownership are similar to those in Cocco (2005), Li and Yao (2007) and others.

Key model parameters are chosen to match moments from the average city based on land scarcity. Cities in the model differ ex-ante only in their land scarcity, and through land scarcity, the stochastic processes for house prices and wages. The model endogenously captures most of the variation by land scarcity in household behavior. Given that, we then perform counterfactual analysis where we vary one element at a time (e.g. vary the level of prices, keeping volatility constant). We find that most of the observed variation in homeownership across cities comes from the observed variation in house price levels and not the variation in risk. However, variation in risk and not prices is important for explaining the variation in LTV.

The weak effect of risk on homeownership is not because the cross-city difference in volatility is small. Cities with land scarcity above the 75th percentile have nearly 50 percent higher standard deviations for each of the innovations to house prices (both those that are correlated and those that

are uncorrelated to innovations in local wages) than cities below the 25th percentile; about the same as the percentage difference in house price levels across these cities. Absent other factors, more risk would lead to more homeownership in the model; homeownership is the only asset available for purchase in the model that has returns correlated with any shock. However, homeownership has many extra costs which potentially increase with more price volatility (such as the transactions costs for buying a house). Moreover, households also have an alternative to using homeownership for insurance: they can accumulate precautionary, non-housing savings instead. We find that, given the extra costs to homeownership, young households with rising income profiles would rather save a little in liquid precautionary savings than save a lot to afford a downpayment. These extra savings help explain the lower LTV ratios in the high risk cities.

In the model, dispersion in price levels has a much larger effect than risk on homeownership choices through the housing ladder. For example, with a housing ladder, a household in our model must own (rent) if it wants to live in a particularly large (small) house. Higher prices in a city decrease housing consumption and can therefore have a large effect on the proportion of households that must own (rent) due to wanting to consume a large (small) amount of housing.

Patterns in the data corroborate our conclusions that differences in price levels cause differences in home ownership rates through housing ladder effects, while differences in LTV are independent of housing ladder effects and are instead due to risk. In the data, once we condition on whether a household lives in an apartment or a house (a proxy for the housing ladder in the model), the negative correlation between homeownership rates and price levels disappear. However, the negative relationship between LTV and volatility does not disappear after conditioning. The substantial variation in homeownership rates across cities is due to the differences in price levels and the presence of a housing ladder.

In the last section of the paper, we discuss why regression-based inferences of risk's effect on homeownership may be biased. Homeownership decisions in economies with transaction costs are durable decisions. Unsurprisingly for an (S,s)-type model, not only contemporaneous prices but also the past history of prices helps determine whether a household currently owns or not. Therefore, homeownership rates within the city economy are also a function of the history of prices. In many studies[6], housing risk is measured using the volatility of area house prices around the time that the homeownership rate is measured. Thus, the volatility variable picks up the history dependence of homeownership on price levels. We show that not accounting for this history dependence in regression-based inferences can lead to biases: specifically, measures of risk are incorrectly found to be important factors in homeownership decisions. The direction of the bias depends on

---

[6]e.g. Sinai and Souleles (2005); Banks et al. (2010)

the realized drift in house prices in the sample, with upward "trends" biasing estimates of the effect of volatility upwards.

A key contribution of this paper is that we provide evidence of a systematic cross-city variation in homeownership and LTV using both micro and aggregate (city-level) data. Banks et al. (2010) uses both variation within and across U.S. states and U.S.-U.K. comparisons on homeownership. Chiuri and Jappelli (2003) looks across developed countries for the effect of financial market imperfections on homeownership. Albouy (2009a) and Albouy (2009b) look at the effects of cross-city variation in taxes and amenities. Han (2010) looks at the effects of housing risks on housing demand and homeowners' propensity to move, using cross-city and time variation. City-level data is appealing since it is more plausible to assume, as we do, that financial market conditions are similar across the areas, in contrast to cross-country comparisons. But, there is still enough plausibly exogenous, observable variation in price levels and risk across cities to find systematic differences in household choices.

Han (2008) builds a model where homeowners may choose to accumulate more housing in order to hedge against housing risks. Under the assumption of separable utility, she provides conditions for when the hedging motive outweighs the household's normal disinclination to hold riskier assets (as in Rosen et al. (1984)). Our work expands on this contribution by adding the option of renting and looking at homeownership and borrowing behavior jointly.

There are a few studies that examine the opposite causal direction - the effect of homeownership and borrowing decisions on prices. Stein (1995) proposes a model where price changes have asymmetric effects on sales due to down payment constraints. Lamont and Stein (1999) finds that cities with high LTVs have higher rather than lower elasticities of house prices with respect to changes in income, but the instrument they use turns out to be weak. Genesove and Mayer (1997) finds that within a specific market (the Boston condominium market), sellers with higher LTVs have higher expected time on the market and receive higher prices.

We do not offer a general equilibrium model of housing, nor do we deal with issues of regional mobility[7] or time-variation in the stochastic process for prices. By using the time-invariant differences in land scarcity across cities to calibrate the different price processes, we largely sidestep issues of endogeneity that may normally arise from examining only one side of a market. Providing structural explanations for the relationship between land values and house prices and for the existence of a housing ladder are interesting explorations that we hope the facts presented here will

---

[7]See Halket and Vasudev (2011) for a model with both, but without the changes in housing supply that we would need here to close our model. Paciorek (2012) has a model of housing supply where the elasticities differ according to factors like land scarcity. Lustig and Van Nieuwerburgh (2010) looks at inter- and intra-regional risk sharing and home values but not homeownership.

encourage.

The rest of this paper is as follows: Section 2 shows the striking variation in homeownership, LTV, house prices and housing risk across U.S. cities, Section 3 presents the model, and Section 4 discusses its parametrization. Section 5 presents our results. Section 6 discusses bias in regressions, and Section 7 concludes. An Online Appendix[8] contains further details, some regression results and robustness exercises.

2   Homeownership and loan-to-value ratios in the data

In this section, we present some basic facts from cross-city variation. First, we show that local homeownership rates are decreasing in price volatility. But, we cannot draw causal inference from this result, because local price volatilities are themselves closely correlated with local price levels. And indeed, homeownership is also strongly negatively correlated with price levels across cities.

We then show that cities with high and volatile prices are also characterized by low LTV ratios. So, households in these cities are less likely to own. And, when they do choose to own, their purchases are less leveraged. But, again, it is not clear whether these outcomes are insurance-type responses to local price volatility, or local prices, or indeed, neither of them.

Local price levels and volatilities are closely related for a simple intuitive reason. They share the same statistical source: variation in land share, i.e. the share of the average city house price that is attributable to the value of land (as opposed to the cost of the structure). As a result, it is not possible to empirically disentangle the impact of price volatilities from that of price levels.

But, we do suggest a suitable instrument for land share, based on local land scarcity. We argue that, controlling for household characteristics (including income), the impact of land scarcity on ownership and LTV comes almost entirely through price levels and volatilities. And so, in the following sections, we simulate a life-cycle model for cities with different values of land scarcity (and therefore, different price levels and volatilities). We find that the model can predict the impact of land scarcity on ownership and LTV reasonably well.

2.1   Data

This study is based on a number of data sources, with our analysis restricted to the cross-section of 2000[9]. Ownership rates and local mean price levels (based on reported values of owned dwellings)

---

[8]available at http://www.halket.com/research/hoinsuranceappendix.pdf

[9]In the Online Appendix, we show that the general patterns also hold for 1990, with the exception of the systematic patterns of LTV discussed below which have become considerably stronger since 1990.

are constructed from the IPUMS 5 percent extract of the US 2000 census, organized by Ruggles et al. (2010). We use two different measures of LTV ratios, taken from the American Housing Survey (AHS) and the Monthly Interest Rate Survey (MIRS)[10]; the latter is maintained by the Federal Housing Finance Agency (FHFA). Quarterly metropolitan house price indices are also taken from the FHFA[11] to estimate local price volatilities. And, our metropolitan-level average earnings series come from the Bureau of Economic Analysis' (BEA) regional program[12]. Finally, we use data on local land scarcity from Saiz (2010) and land share from Davis and Palumbo (2008); we discuss these further below. Where survey data is used, we restrict our sample to households with heads aged 21-75[13], living in houses or apartments.

Throughout, we identify cities with the set of (Primary[14]) Metropolitan Statistical Areas (MSAs), of which there are 297 in the census data in 2000. However, we restrict our sample to the 221 MSAs for which FHFA price data, BEA wage data and the land scarcity instrument are available. Of these, 42 cities are available in the metropolitan surveys of the AHS (for the estimation of local LTV ratios) and in the Davis-Palumbo (2008) data on land shares (itself based on the AHS). And, just 25 are available in the metropolitan MIRS LTV data - though these tend to be the most prominent cities.

Next, we discuss the construction of variables in more detail. First, our measure of local house price volatility is the standard deviation of log annual changes in the FHFA local price index (measured in the first quarter of each year) over the previous five years (1995-2000).[15] This approach is based on Banks et al. (2010). We estimate wage volatility in the same way using the BEA data.

With regards to homeownership, we present much of the evidence on cross-city correlations between ownership rates and price levels/volatilities graphically. But, there are of course concerns that any observed effects will simply be driven by differences in local household composition. Therefore, we construct local ownership rates (from the census extract) that condition on local household characteristics. Specifically, we run a probit-level regression (using the full national sample) of an ownership dummy on various characteristics of the household head[16] together with

---

[10]*http://www.fhfa.gov/Default.aspx?Page=250*

[11]http://www.fhfa.gov/Default.aspx?Page=87. The Case-Shiller indicies distributed by Standard & Poor's (which is the other popular publicly available data set) cover a much smaller sample of cities.

[12]*http://www.bea.gov/iTable/index_regional.cfm*

[13]In the case of the AHS, age is calculated at purchase year.

[14]We do not aggregate these into "consolidated" areas.

[15]The Online Appendix contains robustness results for alternative volatility window lengths.

[16]Quadratic in age, education dummies (high school graduate, 1-3 years of college, 4 year + of college), gender, marital status, dummies for number of children under 18 (1, 2, 3+), ethnicity dummies (black, hispanic) and log household income.

a full set of MSA effects. Then, for each MSA, we predict the ownership rate corresponding to a household with the mean characteristcs in each dimension.[17] To estimate conditional local ownership rates for a particular demographic group (e.g. an age category), we follow exactly the same procedure, but from the beginning (i.e. even before the probit regression) restrict the sample to the relevant demographic group.

As described above, we have two alternative sources of city-specific LTV data: the AHS and MIRS. The AHS is a longitudinal survey, containing detailed information on housing-related vari-ables. The metropolitan survey covers 41 MSAs, and booster samples of a further 6 MSAs (the largest) are included in the national survey[18]. Of these 47 cities, we have land scarcity data on 42; and, we base our AHS analysis on this set of 42. The AHS surveys cover different MSAs in different waves, and we therefore rely on four different waves to put together a complete sample for cross-city analysis: the metropolitan surveys of 1998, 2002 and 2004, and the national survey of 2001. We index observations by year of purchase (rather than survey year), because we have in-formation on the loan amount and home price (to calculate LTV) at the purchase year. We restrict our sample to households with mortgages, and we only study details of mortgages which were taken out when the home was purchased. The last condition ensures that we measure the loan and price in the same year to calculate LTV.

Like with the ownership rates, all reported local LTVs from the AHS are conditional on char-acteristics of the household head[19]. Also, since the AHS samples are not large, we include all households who purchased their home up to five years prior to the survey year (so the full dataset spans purchase years 1993-2004) to predict LTVs in 2000. The consequent overlapping (in terms of purchase year) of the different waves allows us to identify MSA effects. Specifically, we run an OLS regression of LTV on household characteristics, purchase year dummies and MSA dum-mies[20]. And, we predict LTV in each MSA for a household with the mean characteristics in each

---

[17]Because the regression is non-linear, the mean of our "conditional" ownership rates will not equal the mean unconditional rate. Therefore, we recenter all observations by a constant to correct for this.

[18]In the national survey, the samples for cities other than these 6 are insufficient to derive city-specific statistics.

[19]Quadratic in age (at purchase year), education dummies (high school graduate, 1-3 years of college, 4 year + of college), gender, marital status, dummies for number of children under 18 at purchase year (1, 2, 3+), ethnicity dummies (black, hispanic) and log household income (deflated by CPI to 2000 dollars). To predict number of children at purchase year, we count the number of children currently in the family who would have been under 18 at the purchase year; of course, children born between the purchase and survey years are not accounted for.

[20]In this regression, we exclude a number of observations which have suspect LTV values. First, we exclude ob-servations with home purchase prices and loans below $5,000 and LTV ratios above 1.2. Second, the AHS includes a number of imputed values for mortgage size; we exclude these observations, because the imputations are not condi-tional on MSA. There is also a problem with top-coding, discussed in Davis and Palumbo (2008). In the metropolitan surveys, the top code values for house price and loan amount are calculated *by city* (as the mean value of the top-coded observations), which is ideal for our purposes. But, this is not the case for the (relatively expensive) cities with booster samples in the national survey: there, a *national* top code is used. In our sample, 10% of observations in New York

dimension[21], who purchased their home in 2000. To estimate "conditional" LTV for a particular group (e.g. an age category or loan type), we follow exactly the same procedure, but from the beginning (i.e. even before the predicting regression) restrict the sample to the relevant group.

However, there are concerns of large measurement error in the LTV data estimated from the AHS (Lam and Kaul, 2003). And so, we also present our analysis using MIRS data. The MIRS reports (among other statistics) mean LTV ratios for conventional (i.e. excluding federal-guaranteed FHA and VA loans; see Appendix A for description of loan types) single-family loans in 25 cities, based on a monthly survey of mortgage lenders. However, the AHS does have a number of advantages for our purposes: it covers more cities, it covers non-conventional loans also, and (being a household survey) it allows us to control for household characteristics.

Our data on land shares is taken from Davis and Palumbo (2008)[22]. They construct a data set containing, by city and quarter, the average local house price, as well as the share of the local price that is attributable to land value and structure cost, respectively, so that:

$$housevalue_{j,t} = landvalue_{j,t} + structurevalue_{j,t}$$

$$l_{j,t} = \frac{landvalue_{j,t}}{housevalue_{j,t}}$$

where $l_{j,t}$ is then the land share for city j at time t. Their land value estimates are the residual part of house values within a city that are not explained by structure costs. Since it is partially based on the AHS, this data is only available for 42 MSAs in our sample.

As a supply-side instrument for the share of the price attributable to land, we adopt Albert Saiz's (2010) measure of local land scarcity, based on physical constraints on housing supply. For each city, this is the share of a circle around the city center, of 50km radius, that is either steeply inclined land (at an incline of over 15%) or water.[23] Saiz estimates this variable with satellite data.

Throughout this study, we apply city weights to all city-level regressions and scatter plots. These weights correspond to the sum of the households probability weights in each city in the

---

are top-coded, 9% in Los Angeles, and 5% in Chicago. Our approach is to exclude all top-coded observations in the national survey from the regression. However, we take some assurance from the fact that we control in the regression for household characteristics that are correlated with top codes (e.g. household income, education); see the following footnote.

[21]When estimating the means of household characteristics, we include in our sample all excluded observations detailed in the previous footnote. This should partially address the problem of omitted top codes in the national survey (we essentially predict the LTV of the top-coded observations, based on their observed characteristics).

[22]Their data are available at *http://www.lincolninst.edu/subcenters/land-values*.

[23]The data are available at http://real.wharton.upenn.edu/~saiz/SUPPLYDATA.zip. Mian and Sufi (2011) and Chaney et al. (2012) similarly use this data to instrument for elasticities of supply, while Paciorek (2012) builds a model of housing supply that directly connects Saiz's measure of land scarcity to the theoretical elasticity.

census extract, for our chosen sample (households with heads aged 21-75, living in houses or flats). And, all statistics from the census extract and AHS are estimated using the available household probability weights.

2.2 Homeownership and LTV

As the first panel of Figure 1 shows, homeownership is negatively correlated across cities with house price volatility. The predicted (OLS) effect shows ownership rates ranging from about 0.7 (for the least volatile cities) to 0.4 (for the most), with an R squared of 41 percent. But, as the second panel shows, it is also strongly negatively correlated with price levels (see the first two panels of Figure 1): here, the correlation is 60 percent. It should be noted that New York appears to be an outlier in these homeownership figures. However, we propose an explanation in Section 5.

Unsurprisingly, volatilities and levels are themselves closely correlated, with a correlation of 38 percent (see Figure 2). Consequently, it is difficult to disentangle their respective effects. To see this graphically, we isolate the portion of variation in volatilities that is uncorrelated with price levels (i.e. the residuals from an OLS regression of volatilities on levels). In the first panel of Figure 3, we plot homeownership against these volatility residuals: the effect is much weaker than before, with less than half the coefficient and an R squared of under 5 percent. In the second panel, we plot homeownership against price level residuals (from a regression on volatility). The relationship is stronger than the one from the volatility residuals[24], though still much weaker than in Figure 1: the correlation is 23 percent.

As with homeownership, average city loan-to-value ratio at origination (LTV), as measured by the AHS (controlling for household characteristics), is strongly negatively correlated with both local price volatility and level (Figure 4). In the first panel, moving from the lowest to highest price volatilities in the MSA sample, following the OLS-predicted line, LTV falls from 0.85 to 0.79, with an R squared of 19 percent. But again, unsurprisingly, in the second panel, there is also a strong negative correlation between LTV and price levels (R squared is 50 percent).

These relationships also exist for the MIRS data, as Figure 5 shows. The correlation is much tighter than for the AHS data: the AHS is likely to be subject to substantial sampling and reporting error. But, the magnitudes of the effects on LTV are very similar for the two datasets: for example, for volatility, the MIRS effect is -1.2 compared to -1.0 for the AHS. Notice that the mean LTV is lower for the MIRS (0.77) than for the AHS (0.84) estimates. This is in part because the MIRS sample is restricted to conventional loans only: the AHS mean for conventional loans is 0.80.

---

[24]This is perhaps unsurprising given that volatility is measured with more error than levels.

It might be argued that these LTV patterns have a supply-side explanation, due to the intricacies of American mortgage institutions. But, in Appendix A, we show that non-varying conforming loan limits do not drive the observed cross-city variation in LTV, and nor do differences in local mortgage interest rates. Also, there may be concern that the variation in LTV is merely arising from cross-city differences in shares of mortgage-holders. Indeed, almost a quarter of homeowners (in the 5 percent census extract of 2000) do not hold mortgages. However, it turns out that the local mortgage share (among homeowners) is uncorrelated with the ownership rate itself, so it is not likely to be driving our results.

To summarize, expensive and price-volatile cities tend to be characterized by low ownership rates, but also low LTV ratios. Households in these cities are less likely to own, and when they do buy a house, they take a bigger equity stake. But, we cannot make causal statements based on this evidence, given the close association between local price volatility and levels.

2.3   Association between house price volatility and levels

To understand the close link between volatilities and levels, it is necessary to view house prices as the sum of two components: land values and structure costs. Price volatilities and levels are correlated across cities because they share the same statistical source: variation in local land shares. In this subsection, to ensure comparability, all data (including house prices) are taken from Davis and Palumbo (2008) - for 42 MSAs.[25]

Consider first the cross-city variation in price levels. The first panel of Figure 6 shows that there is substantial variation across cities in house price levels (the range covers two log points). But, comparing the final two panels of Figure 6, the cross-city variation in structure costs is negligible: it is land values that are driving the large variation in house prices. Clearly then, cities with larger land shares will have higher house prices.

Next, consider the variation in house price volatilities, i.e. the standard deviations over annual growth rates. As with the price levels, the first panel of Figure 7 reveals large variation in volatilities, ranging from 0.006 to 0.075. Since the covariance between the growth rates of land and structure costs within a city over time is, on average, negligible, the standard deviation of house price growth over time within a city can be approximated as follows:

$$\sigma_j(g_{jt}^{hp}) \approx l_{j,t}\sigma_j(g_{jt}^{lv}) + (1 - l_{j,t})\sigma_j(g_{jt}^{sc}), \tag{1}$$

---

[25]It should be noted that the mean house price level is 16% larger in this data than our larger census-based sample, and the mean local price volatility is also slightly larger (0.030 compared to 0.025 in the FHFA). But, the Davis and Palumbo dataset allows us to examine the disaggregated components of local house prices.

where g_{jt}^{lv}, g_{jt}^{sc} and g_{jt}^{hp} are the annual growth rates of land values, structure costs and house prices, respectively. Empirically, house price volatility, $\sigma_j(g_{jt}^{hp})$, is uncorrelated with the volatilities of local land values, $\sigma_j(g_{jt}^{lv})$, and structure costs, $\sigma_j(g_{jt}^{sc})$. But, critically, Figure 7 shows that the variation in land value is an order of magnitude larger than the variation in structure costs. Therefore, according to Equation 1, house price volatility should be strongly positively correlated with land share - through a composition effect.

And indeed, Figure 8 confirms that both local house prices levels and volatilities are increasing in land share[26]. In further results in the Online Appendix, we confirm that the effect on price levels comes entirely through the land value component (structure costs are uncorrelated with land share). And, the observed relationship between land share and local price volatility is entirely a composition effect: land values are much more volatile than structure costs (in fact, land value and structure cost volatilities are individually uncorrelated with land share).

2.4   An instrument for land share

Given the strong empirical connection between price volatility and levels, we have to disentangle their effects on household choices with a model. One approach would be to simulate ownership and LTV decisions in cities with different land shares. These cities are characterized by different local price volatilities and levels (this is what matters for the simulation), and these can be estimated from the data (by reference to their empirical relationship with land share).

The problem is that the correlation between land share and ownership or LTV cannot be considered causal *a priori*. Omitted variation in local productivity or housing demand should not be a concern, because we are controlling for household income in our estimates of ownership and LTV. But, we are worried about reverse causation from ownership/LTV to price levels/volatilities and land share.

There are two ways to address this issue. The first is to simulate a general equilibrium model, where land share, price levels and volatilities are all determined endogenously. However, the determination of these housing market outcomes is not the focus of this paper, and the equilibrium conditions would significantly complicate computation.

The second approach, which we choose, is to find a suitable instrument for land share that will only affect ownership and LTV indirectly, i.e. via local price levels and volatilities. We opt for a measure of local land scarcity, described in the data section above. It is based on local geographical features, namely inclined land and water. Conditional on household income, this instrument is unlikely to affect tenure and LTV in a significant way *directly*; the effect should only

---

[26]Reported land share is the mean over the 4 quarters of 2000, based on estimates from Davis and Palumbo (2008).

come through local housing conditions (captured by price levels and volatility). The first stage is sufficiently powerful. Figure 9 shows a strong relationship between land scarcity and land share: a 1 percentage point increase in land scarcity is associated with a a 0.5 percentage point increase in land share. Unsurprisingly, Figure 10 shows there are strong positive relationship between land scarcity and price volatilities/levels as well, with correlations of 28 and 30 percent respectively.

As above, we confirm in the Online Appendix that the entire effect of land scarcity on price levels comes through the land value component (and not structure costs). And, the effect on price volatility is entirely a composition effect: land scarce cities have larger land shares, and local land values are more volatile than structure costs.

3 Household choice model

In this section, we build a life-cycle model of households that work and consume in a particular city for their entire lives. Several of the assumptions we make deserve extra attention. We severely limit households' access to insurance in a way which should bias the model *in favor* of using homeownership as insurance: we do not allow for inter-city migration, so households cannot use moving away from the city as source of insurance against house price changes; and the only asset besides a house is a risk-free bond. Also, though the model is "partial-equilibrium," rental prices are tied to sale prices through an implied equilibrium relationship that leads to counterfactually high rental volatility. In the model, rents will be as volatile as house prices, while in the data they are clearly lower[27]. Excess volatility in rents as compared to prices will again bias the model in favor of homeownership as insurance.

Time is discrete, and each period in the economy corresponds to one year in the data. Households are born at age $a = 21$ and live at most to age $a = 75$. A household is indexed by i and lives in a city, indexed by j, for its entire life. The city has a time-invariant land scarcity λ_j.

3.1 Preferences

Households have recursive preferences of the Kreps and Porteus (1978) type[28]. The household gets instantaneous utility from a non-durable consumption good c and a durable housing good h according to:

$$u(c_t, h_t, a_t) = (c_t^{1-\sigma} h_t^{\sigma})/F(a_t)$$

---

[27]For instance, see Campbell et al. (2009) and Verbrugge (2008).

[28]These preferences nest time-separable preferences but allow for the separate consideration of inter-temporal smoothing (savings) and smoothing across states within a given period (risk-aversion).

The path for the family size adjustment factor, $F : \{21, 22, ..., 75\} \rightarrow \mathbb{R}_{++}$, is exogenous, constant across households of the same age and known to the household at birth.[29] The household's utility at time t, V_t, is then given by the composite of its instantaneous utility and its future expected utility:

$$V_t = [(1-\beta)u(c_t, h_t, a_t)^{1-\phi} + \beta(\mathscr{R}_t V_{t+1})^{1-\phi}]^{\frac{1}{(1-\phi)}},$$

where future expected utility is given by $\mathscr{R}_t V_{t+1} = (\mathbb{E}_t[V_{t+1}^{1-\gamma}])^{\frac{1}{1-\gamma}}$. γ measures risk aversion while ϕ is the inverse of the intertemporal elasticity of substitution. Additive utility is a special case where $\phi = \gamma$.

Households get utility at death from bequeathing wealth, $V_{t+1}(\cdot, a_t = 75) = (b_{t+1} + p_{t+1}^j h_t)^{1-\sigma}$.

3.2   Labor Earnings

Households receive labor earnings, Y_t[30] up until an exogenously set retirement age R, after which they receive a pension. Y_t contains three components: a deterministic life-cycle component, an idiosyncratic component and a city-specific component. The city-specific component W_t^j, which we call wages, follows a geometric random walk. The idiosyncratic component is a geometric random walk with a transitory shock:

$$Y_t = L_t^i W_t^j \rho_t^i$$
$$L_t^i = \exp f(a_t) L_{t-1}^i \psi_t^i$$
$$W_t^j = W_{t-1}^j v_t^j$$

where $\ln \psi_t^i \sim \mathcal{N}(-0.5\sigma_\psi^2, \sigma_\psi^2)$, $\ln \rho_t^i \sim \mathcal{N}(-0.5\sigma_\rho^2, \sigma_\rho^2)$ and $\ln v_t^j \sim \mathcal{N}(\mu_v - 0.5\sigma_v^2(\lambda_j), \sigma_v^2(\lambda_j))$. The variance of innovations to wages, v_t^j, can differ across cities according to their land scarcity; however all cities have common drifts. After retirement, the household gets a proportion (adjusted

---

[29]Attanasio et al. (1999); Gourinchas and Parker (2002); Cagetti (2003); Li and Yao (2007) each let family size affect a household's discount factor. In Gourinchas and Parker (2002); Li and Yao (2007), the life cycle profile for family size is deterministic and homogeneous across households of the same age. Attanasio et al. (1999); Cagetti (2003) let the profiles vary by education. Browning and Lusardi (1996) have a stochastic process for family size (see their paper for more references). Gervais (2002); Campbell and Cocco (2003); Li and Yao (2007); Diaz and Luengo-Prado (2008) all use Cobb-Douglas preferences over non-durable consumption and housing which are consistent with evidence from Davis and Ortalo-Magne (2011) that housing expenditure shares are approximately constant across cities.

[30]For ease of notation, we suppress some variables' dependence on the household and city specific labels, i and j respectively, when such dependence is or will become obvious.

for growth in the city) of its final salary, $Y_t = \zeta L_R^i W_t^j$. All households' income is taxed at a rate t_y.

3.3 Housing Market

At any time, homes may either be rented ($\tau_t^i = 0$) or owned ($\tau_t^i = 1$), but not both simultaneously. There is a housing ladder which forces households to choose rented housing from the set H^r and owner-occupied housing from the set H^o.

Housing can be bought at a unit price p_t^j, which contains two components - one correlated with labor earnings and one uncorrelated with labor earnings:

$$p_t^j = Q_t^j W_t^j$$

where $Q_t^j = Q_{t-1}^j \varepsilon_t^j$ and $\ln \varepsilon_t^j \sim \mathcal{N}(-0.5\sigma_\varepsilon^2(\lambda_j), \sigma_\varepsilon^2(\lambda_j))$. The variance of innovations to the uncorrelated component, ε_t^j, like those of the correlated component, differs across cities; city-specific drifts remain common[31].

An owner pays proportions, t_p and δ_j, of the value of the house each period towards property taxes and maintenance, respectively. The housing maintenance means houses do not depreciate and the maintenance required may vary across cities. A household may not "build on" to its house; to adjust the size of an owner-occupied house, it must sell its current one and buy a new house. Each time a household buys a house, it pays a fraction θ_b of the value of the house as a transaction cost.

A renter pays only the spot rental price per unit of housing s_t^j, which we set so that a risk-neutral landlord would be indifferent between renting or selling the house, subject to paying income tax on its rental income:

$$s_t^j = \frac{t_p + \delta_j + \frac{r_b - \mu_v}{1 + r_b}}{1 - t_y} p_t^j$$

where r_b is the risk-free interest rate at which households and landlords can borrow.

Households have three potential motives for owning: renting is higher than the user-cost of owning due to the taxation of rental income, the housing ladder restricts the size of rental housing, and the several insurance motives.

---

[31]Flavin and Yamashita (2002); Campbell and Cocco (2003); Cocco (2005); Yao and Zhang (2005); Li and Yao (2007); Diaz and Luengo-Prado (2008) all assume shocks to house prices are permanent.

3.4 Assets

Besides housing, the only other financial asset for the household is a risk-free one period bond, b_{t+1}^i, which pays r_l to savers but costs (net) $r_b > r_l$ to borrow. Households may borrow at this rate, subject to a borrowing constraint. Housing is the sole form of collateral. We model this by giving households a home equity line of credit.[32] The LTV at the time of purchase is simply the ratio: $-\frac{b_t^i}{p_t^j h_t^i}$.

When purchasing a home, households can borrow up to $(1 - d)$ of the value of the house, where d is the down payment constraint. Thereafter, as long as they continue to be homeowners, agents may borrow up to $(1 - d)$ of the value of the house. They may also choose to roll over their debt after making an interest payment. So at any time, the borrowing constraint is:

$$b_{t+1}^i \geq \min\{-(1 - d)\tau_t^i p_t^j h_t^i, (1 - 1_m)b_t^i\},$$

where 1_m is an indicator variable which equals one if the household chooses to move in the period.[33]

If the household chooses to sell its home, it must pay off all existing debt, though another loan can be taken out if another home is purchased. A household that does not have positive total cash-in-hand (housing wealth plus financial wealth plus current income) will not be able to pay off the mortgage it has (the debt it owes) on its home and will not choose to move in this period. We do not allow the household to choose to default (see Jeske and Krueger (2005) for a model with mortgage default), but households can default implicitly by dying. After retirement, we do not allow households to take out new loans, but they may continue with their old loan[34]. This effectively ensures in our calibrated economy that all households reach age 75 debt free.

Newborn households are "born" with no housing but they draw their initial wealth from a distribution Π_b, which is a probability distribution on $\mathbb{R}_+$.

---

[32]We also call this a mortgage throughout.

[33]This borrowing constraint is different from the more typical one which restricts borrowing to be weakly less than some percentage of the house value ($b_{t+1}^i \geq -(1 - d)\tau_t^i p_t^j h_t^i$). With risky house prices, for an agent near the typical borrowing constraint, a fall in the value of a house results in a "call on the mortgage principle" - the agent must reduce the amount borrowed. If house price volatility is large enough, the effective down payment constraint (the amount the agent could borrow and still be able to repay in any state of the world next period) may be much tighter than the actual (d).

[34]That is, a retired household's borrowing constraint is $b_{t+1}^i \geq \min\{0, (1 - 1_m)b_t^i\}$.

3.5 Household's Problem

The problem of the household is to choose consumption, house size and ownership, and savings, given its permanent and transitory earnings components, housing and assets at the beginning of the period and prices, subject to budget, borrowing, and choice-set constraints and the initial condition and laws of motion for Q_t^j, W_t^j (which we do not repeat below) for all variables:

$$V(a_t, L_t, \rho_t, b_t, \tau_{t-1}h_{t-1}; Q_t^j, W_t^j, \lambda_j) = \max_{c_t, h_t, b_{t+1}, \tau_t} [(1-\beta)u(c_t, h_t, a_t)^{1-\phi}$$

$$+ \beta(\mathscr{R}V(a_{t+1}, L_{t+1}, \rho_{t+1}, b_{t+1}, \tau_t h_t; Q_{t+1}^j, W_{t+1}^j, \lambda_j))^{1-\phi}]^{\frac{1}{(1-\phi)}}$$

s.t.

$$c_t + b_{t+1} + h_t((1-\tau_t)s_t^j + \tau_t p_t^j(\delta_j + t_p + 1 + 1_m\theta_b)) \leq b_t(1+r) + Y_t(1-t_y) + h_{t-1}\tau_{t-1}p_t^j$$

$$b_{t+1}^i \geq \min\{-(1-d)\tau_t^i p_t^j h_t^i, (1-1_m)b_t^i\}$$

$$r = \begin{cases} r_l & \text{if } b_t \geq 0 \\ r_b & \text{if } b_t < 0 \end{cases}$$

$$c \geq 0 \qquad \tau_t h_t \in \{0, H^o\} \qquad (1-\tau_t)h_t \in \{0, H^r\} \qquad \tau_t \in \{0,1\}$$

4 Parametrization

Parameters that vary across cities in the model are those indexed by j and differ ex ante according to their land scarcity, λ_j. All other parameters remain constant across cities. In this section, we discuss the calibration/estimation of some key parameters; the calibration of the remainder are discussed in the Appendix (see Table 1 for their values). These key parameters are all those that vary across cities and the housing ladder parameters in H^r and H^o. These are estimated in three steps.

1. We initialize the model so that the cross-section of relative prices and wages in a particular year, 2000, is the same in the model as in the data. We assume $\sigma_v^j, \sigma_\varepsilon^j, p_{2000}^j, W_{2000}^j, \delta^j$ vary

across cities in the model according to the same (linear) relationship estimated in the data.

$$\sigma_v^j = \alpha_v + \beta_v \lambda_j \tag{2}$$

$$\sigma_\varepsilon^j = \alpha_\varepsilon + \beta_\varepsilon \lambda_j \tag{3}$$

$$p_{2000}^j = \alpha_p + \beta_p \lambda_j \tag{4}$$

$$W_{2000}^j = \alpha_w + \beta_w \lambda_j \tag{5}$$

$$\delta^j = (1 - \alpha_\delta - \beta_\delta \lambda_j) \delta_h \tag{6}$$

We estimate the parameters of these functions using land scarcity, as discussed in Section 2 (see below for more detail). Further details about the estimates of the α and β coefficients can be found in the Online Appendix. The δ^j are set using the relationship between land share and maintenance described below. This entire step can be done without computing the household's problem.

2. We simulate a set of cities with different land scarcities, each with 200,000 households. Each household is born at some date at most 54 years before 2000. For each city, we draw realizations of the annual innovations to prices and wages so that they equal their 2000 relative value in 2000.

3. We choose the parameters in the housing ladder so that specific moments in the simulated model data best match those in the data in 2000. The values of the housing ladder parameters are found by repeatedly computing the household's problem for different values of the parameters (and repeating step 2) and choosing the pair that provided the best match.[35]

4.1 Housing

We assume that a city's housing supply is fixed and that homeowners pay a maintenance cost to replace depreciated housing capital. So, the (growth-adjusted) relationship between housing depreciation and housing investment is (abusing notation)

$$\delta_h = \frac{I_h - \Delta(pH)}{pH}$$

For the aggregate value of housing, pH, we use non-farm owner-occupied housing from NIPA's *Historical-Cost Net Stock of Residual Fixed Assets* table. Investment in housing is computed using

---

[35] A simulated method-of-moments computed over a grid of potential parameter values.

non-farm owner-occupied housing from NIPA's *Historical-cost Investment in Residential Fixed Assets.* This gives $\delta_h = 0.017$. These values from NIPA are the value of the structures and do not include the value of land. For any city, j, $1 - \alpha_\delta - \beta_\delta \lambda_j$ is the share of structure costs in house value. So for each city, we set $\delta_j = (1 - \alpha_\delta - \beta_\delta \lambda_j)\delta_h$, where $\alpha_\delta = 0.306$ and $\beta_\delta = .470$ are the intercept and slope, respectively, of the linear relationship estimated between land share and land scarcity (see the Online Appendix for further details). The rent-to-price ratio in the cities will therefore vary slightly with land scarcity due to changes in δ_j.

We allow households to choose any size rental up to a maximum: $H^r = (0, \overline{h^r}]$. We impose a minimum owner-occupied house size but no other restriction: $H^o = [\underline{h}^o, \infty)$. We use the model to set $\underline{h}^o$ and $\overline{h^r}$ so that cities with the mean land scarcity have an average homeownership rate in the model matches the fitted homeownership rate at the mean land scarcity in the data and so that the mean ratio of owner-occupied house sizes to rental house sizes in the model matches the fitted ratio (in square feet) in the data. Matching the two moments, the homeownership rate and the relative housing sizes, identifies the two parameters uniquely. We do not have a formal proof but casual introspection (if $\underline{h}^o$ increases, then $\overline{h^r}$ must decrease to keep the ownership rate constant, but $\overline{h^r}$ must increase to keep the relative house size ratio constant) and all computation thus far confirms it.

4.2 Prices

To estimate the parameters in the price processes, we match year/city panels of house prices (from the FHFA) and average wages (from the BEA). These data are used to calculate, for each city, a covariance matrix of annual growth rates of wages and house prices over 1995-2000. We have assumed in the model that v_t^j affects prices and wages equally. So we could use either wage growth variance or price-wage growth covariance as alternative estimates for $\sigma_\varepsilon^2(\lambda_j)$. The mean (across cities) wage growth variance (0.00021) is almost twice as large as the mean price-wage growth covariance (0.00011). However, if we restrict our attention to the thirty largest cities in the sample, the two statistics do match (they are both 0.00015). So we choose to use the wage growth variances to estimate α_v and β_v.

We run OLS regressions on equations (2) to (5), with land scarcity as the independent variable, and a range of dependent variables: local price volatility (the standard deviation over annual growth rates, 1995-2000), wage volatility, price level and wage level. We take local wage levels from the BEA data of 2000, and estimate price levels (as discussed above) from the 5% census extract of 2000.

Conditional on α_v and β_v, house prices are used to estimate α_ε, β_ε and β_p. Due to the homo-

geneity in our model and since we only set the housing ladder parameters in a later step, we are free to normalize α_w and α_p. Table 2 shows some moments for the key parameters. The results from the instrumental variable regressions are available in the Online Appendix.

5 Results

5.1 Moments in models and data

Table 3 shows the results from the average city by land scarcity compared to the data. Since the house size parameters were chosen so that the model matched the data on the homeownership rate and relative house sizes, it should not be surprising that we attain a very good fit along these lines. The model also matches the data well if we consider only those households 65 years old and younger which, given the model's relatively simple characterization of post-retirement life, is also not surprising. Table 4 shows that the model also matches the profile of homeownership relatively well, though there are too few young and too many middle-aged homeowners.

Though no parameters were chosen to match the LTV rates (conditional on taking a loan), the model is able to match the data from the AHS relatively well, however it is somewhat lower. This is perhaps a result of only having one non-housing asset in the model. In the data, we do not observe the mortgage net of other financial assets, which is the relevant variable in the model.

Table 5 shows the slopes of linear regressions of city-level homeownership and LTV on land scarcity in the data and in the model. The model is able to explain much of the difference in homeownership and LTV across cities. A ten percentage point increase in land scarcity implies a decrease in homeownership of 2.5 and 2.0 percentage points in the data and model, respectively. Likewise, the same increase in land scarcity implies a decrease in LTV of 0.7 and 0.6 percentage points in the data and model, respectively. Generally speaking, the difference in homeownership rates across land scarcity declines with age, a pattern which the model matches.[36]

The model also explains the difference in LTV over cities for the younger age groups (matching well up to age 45) but breaks down thereafter. As in the data, the relationship between LTV and land scarcity becomes less negative with age, turning positive for the older ages (although the coeffecients are not significant for these ages). The shortcoming of not being able to observe net financial assets in the data is likely to be more acute for older households that have accrued savings, and perhaps explains why the increase in the coefficients is sharper in the model than in the data. Fortunately, late-life LTV figures are relatively inconsequential for the cross-city dispersion: in

---

[36]This may explain why Sinai and Souleles (2005) finds a positive effect on ownership from the interaction between mobility (which is negatively correlated with age) and rent volatility (which is positively correlated with land scarcity).

the data, 80 percent of new loans are taken by households under 50 and 97 percent by households under 65. Thus the restriction that, in the model, households over 65 are not allowed to take new loans is probably not important for the LTV results.

5.2   Counterfactuals

High land scarcity cities differ exogenously in several important ways from low land scarcity cities in the model. High land scarcity cities have lower rent-to-price ratios in the model because structural depreciation as a percentage of the house value is smaller (this is also true in the data). High land scarcity cities have higher idiosyncratic house price variance and a higher variance in the innovations to the correlated house price - wages process. These cities also have higher prices but also, as a mitigating factor, higher wages.

To disentangle these different contributions, we simulate five variations to the baseline model economy, each time allowing only one of the parameters to vary with land scarcity; all other parameters are kept at their respective mean land scarcity values. In two of the counterfactuals, we allow the variances to vary by land scarcity according to equations 2 and 3, respectively. In a third and fourth, we simulate cities with different land scarcities so that they have relative prices or wages in the year 2000 that vary according to equations 4 and 5, respectively. In the final counterfactual, we vary the maintenance in cities according to equation 6. Table 6 shows the coefficient from regressing homeownership and LTV on land scarcity from each of the counterfactual economies, so highlighting each parameter's contribution to the cross-city differences generated by the model.

The largest contributor to the cross-city dispersion in homeownership is dispersion in the level of house prices. However, higher risk explains why high land scarcity cities have lower LTVs. Changes in risk do affect homeownership slightly. But the results show that higher risk reduces homeownership; households, on balance, do not use homeownership to insure themselves against housing risk.

5.2.1   Differences in price levels and homeownership

Differences in prices create differences in homeownership through the housing ladder. In both the data and the model, households live in larger houses when they live in cheaper cities, and the difference in sizes is larger for owners than for renters[37]. So, households in cheaper cities, living

---

[37]The square-footage difference in the data (using the AHS, all reported differences statistically significant) between the 75th percentile and the 25th percentile city by land scarcity is 21 percent, while in the model the size difference is 19 percent of the average size house. For owners, the difference in house size is 15 percent and 17 percent in the data and model respectively. For renters the difference is 13 percent and 9 percent, respectively.

in larger houses, are more likely to choose to own due to a binding maximum rental constraint, while households in the expensive cities are more likely to rent due to a binding minimum owner-occupied house size constraint. Likewise, differences in wages work similarly, though the total effect is smaller. Everything else equal, higher wages in the high land scarcity cities leads to higher housing consumption and, due to the housing ladder, higher homeownership (consistent with findings from Coulson and Fisher, 2009).

In our simulations, households respond to lower prices by increasing housing consumption, which leads to higher homeownership due to the housing ladder. Therefore, households who adjust their tenure decision in response to local prices are likely doing so to adjust their housing consumption. And so, holding housing consumption fixed, we should expect to see no effect of price on tenure decisions.

In the data, a critical margin of adjustment in housing consumption is between apartments and houses.[38]According to our census sample, almost all houses (85 percent) are owned and almost all apartments (87 percent) are rented. Interestingly though, among owned properties, LTV ratios (predicted for 2000 from the AHS) are almost identical across dwelling types: 0.84 for houses and 0.86 for apartments.[39]

In Figure 13, we plot the relationship between conditional homeownership rates and prices across MSAs, by dwelling type: full sample, houses only and apartments only. The estimated effect of prices on ownership rates is more than three times as large for the full sample as the houses-only sample. And, the relationship for apartments is actually positive. Clearly then, adjustments in housing consumption (in this case, between dwelling types) play an important part in driving the overall price-ownership relationship. Also, Figure 13 suggests that New York is an outlier in Figure 1 because it has more apartments per unit of housing than the typical city with its land scarcity.

In contrast to the tenure choice, the LTV decision (conditional on ownership) is not strongly related to housing consumption. This suggests the mechanism driving these LTV results is independent of the housing ladder (instead, we argue below that insurance motives are important). In Figure 14, we plot LTV-price relationships for the full sample, and separately for houses and apart-

---

[38]Based on US census data, we have defined an "apartment" as a housing unit that shares its structure with one or more other housing units; a "house" is a single-unit structure. Note that "houses" need not be entirely detached from other structures: a housing unit attached to another unit by a full-height dividing wall, that goes from ground to roof, is here defined as a "house".

[39]These two statistics are means of local conditional LTVs across the 42 cities in our sample (for houses and apartments respectively), where the conditional LTVs are estimated as described in Section 2.1. Controlling for household characteristics (as we do in this procedure) is not trivial because of the omitted top-coded observations (see Section 2.1 for details).

ments. This time, the effects of log prices are negative for *both* dwelling types; they are also very similar in magntiude: -0.060 for houses and -0.064 for apartments. It is clear that composition effects are not driving the relationship for LTV.

5.2.2 Differences in risk and LTV

Higher risk has a small effect on homeownership, but accounts for all of high land scarcity cities' lower LTV. This is for two reasons. Firstly, households save more in the high variance cities, in part due to higher price volatility but also due to high wage volatility. Renting is a partial hedge against falls in wages that are correlated with rents and prices. However, households will not completely insure themselves against falls in wages through rental housing, since doing so would distort their housing consumption greatly. So households also hold more total wealth in the high co-variance economies. This extra total wealth leads to lower LTVs when the households do decide to purchase a house.

Model households do view homeownership as a potential source of insurance; but it is highly imperfect variety of insurance. Housing comprises about 25 to 30 percent of expenditures in the model, so households would like homeownership to comprise about the same amount in their total wealth portfolio (including human capital) for insurance purposes. However, most renters would have to leverage their financial wealth greatly to buy a home, leaving them near the borrowing constraint and particularly exposed to large falls in house prices, which are more likely in risky cities. Large falls in prices can leave their budget sets particularly small. So households in riskier cities defer housing purchases until they can afford to buy the house with lower leverage.[40] This is consistent with our finding that the share of homeowners that purchase their house without a mortgage is not correlated with risk (or land scarcity) in the data: households with this much financial wealth in the model are not troubled by the borrowing constraint.

There are also transactions costs: more risk leads to more mobility and lower expected durations in any given house. Since adjusting owner-occupied housing is costly, this decreases the value to owning and thus homeownership. For instance, in the σ_ε counterfactual, the homeownership rate for households under 35 in a typical low land scarcity city (25th percentile) economy is 2 percentage points higher than its counterpart at the 75th percentile. If we eliminate transactions costs, the (negative) slope of homeownership with respect to land scarcity halves in both risk counterfactuals. In other words, households optimally prefer to self-insure with a risk-free bond, which does not have transactions costs, does not distort the intratemporal consumption bundle and

---

[40]These are two sides of similar coins: households worry about the risk that their house will be expensive at the time of purchase and also the risk that their house falls in value after buying it.

does not compel asset-poor households (which young would-be homeowners largely are) to over-leverage themselves, rather than insure with housing even though housing is the only asset whose return is correlated with the some of the risks the household faces.

Households do own slightly larger houses in the higher variance economies but there is little evidence of a housing ladder effect as discussed in Banks et al. (2010). Their theory is that households that expect to consume more housing than a rental can provide later in life anticipate owning later in life and therefore households in economies with high risk will seek to insure themselves against the risk that prices may be high in the future, when they are likely to own a large house, by purchasing a small house earlier.

Theoretically, the overall strength of the "ladder effect" is particularly dependent on the nature of the housing ladder assumptions. A very rigid ladder (where, say, the minimum owner-occupied size equaled the maximum rental available, such as in Banks et al. (2010)) can potentially have large average effects early in the life cycle. However, a rigid ladder with a low maximum rental size would not enable our model to match the relative housing consumption of renting versus owning households seen in the data. More importantly, if the ladder effect were large, new homeowners should be willing to buy housing with lower down payments in order to own sooner in riskier cities. From the price level counterfactual, we do see that households would opt for higher leverage purchases in expensive cities in order to climb the housing ladder. If more risk also led households try and climb the ladder faster via larger loans, the model would not be able to match the higher down payments (lower LTVs) in land scarce cities in the data.

Finally, the effect of differing maintenance costs is small: maintenance after all is only part of the cost of housing. Relatively high maintenance in low land scarcity economies makes owner-occupancy relatively more attractive as it increases the tax wedge in the user-cost formula, leading to very slightly higher homeownership rates and LTV ratios.

6   History dependence and regression-based inference

A household's decision to become a homeowner is a durable decision. The durability of this decision means that at the aggregate (city) level, homeownership rates are not only a function of contemporaneous prices but also of lagged prices. Not accounting for this history dependence in regression-based inference can lead to biases which make measures of volatility appear as important factors in homeownership.

6.1   The history dependence of homeownership rates

We can use the model to illustrate this history dependence, though the point carries in any setting where housing is a durable choice (for instance, whenever there are transaction costs). We preform the following experiment. We take four cities (two sets of two) that are each ex-ante identical (all four have the have the mean value of land scarcity). In the first set, we give one city a positive price shock in period one and a twin city the same positive price shock in period five[41], so that prices are identical from year five onwards. In the second set, we apply the same magnitude and timing of shocks, though negative rather than positive.

Figure 15 plots the price path for each city, the implied measured volatility using a 5 year rolling window[42]. and the path of homeownership in each city. This figure shows two important characteristics. First, homeownership rates do not instantaneously fully incorporate the changes in prices (a "delay effect"). Second, there is an asymmetry in the speed with which a city adjusts to shocks. Positive shocks feed into homeownership rates faster than negative shocks.

The model is homothetic, so a positive shock to prices and wages is equivalent to a shock which proportionally reduces (towards zero) all households' savings or borrowings. Renters with a lot of financial wealth feel "wealthier" ("poorer") after a negative (positive) shock to prices and wages and so expect to consume more (less) non-durable and housing goods over the rest of their lives. In particular, they may now expect to cross (or not cross) one of the housing consumption boundaries over their life cycle. In general, households want to turn this new-found wealth into consumption over the rest of their lives. The pivotal change in housing consumption (with respect to the housing ladder constraints) may not be until later in their lives. So homeownership rates take some time to fully react to the change in prices. Wages also go down (up) though, so young, renting households with little financial wealth who are years away from becoming owners are hardly effected by the shock. Eventually, as existing households exit the model and newborn households enter, the ownership rate begins to revert back to its steady-state level.

The shocks have asymmetric effects due to the transactions cost. A negative (positive) shock expands (contracts) the contemporaneous consumption set of the household. In particular, a negative shock means that a household that wanted to presently own but could not due to the borrowing constraint might now be able to (and vice versa). Most households in such a position are young households, as they have much less wealth on average. But younger households also expect to move more often than older households due to higher earnings uncertainty and lower wealth. So

---

[41]In each case, we use the shock that is correlated with wages, v_t^j.

[42]Sinai and Souleles, 2005 and Banks et al., 2010 use a nine-year window and a five-year window, respectively, in their regressions.

they find the transactions cost of buying a house more onerous. So whereas an increase in house prices will strictly prevent some households from becoming a homeowner in the near future and therefore has a more immediate effect on homeownership rates, a decrease in prices merely allows some households to choose to become owners - a choice they may defer for a while due to mobility expectations. Put another way: an increase in prices may bound some households away from owning which means they *must* rent, whereas a decrease can slacken the borrowing constraint for some households which means they *can* own.

6.2   Regression-based inference

The delay effect and the asymmetric effect have important implications for inferences from within-city variation (across time) in price volatility. The delay effect means that cities with the same prices and same stochastic processes may still have different homeownership rates because their price histories differ. This can be clearly seen in the first panel of Figure 15, where after year five, the homeownership rates differ across cities that received the same shock at different points in time. The asymmetric effect means that if one were to look at absolute price changes (or squared price changes) the delay effect would not average out.

There is obviously measurement error in using realized volatility measured from recent historical prices as the proxy for ex-ante risk. From the durability and asymmetry effects we know that (1) the history of prices is an important explanatory variable for homeownership rates and that (2) volatility measures, which use squared historical prices, will be correlated with contemporaneous homeownership rates, even if in the true DGP risk is not. In other words, even if the latent variable, risk, is not important for homeownership rates, the measurement error part of volatility is important.

Lastly, the amount of bias that arises from using historical volatility as a measure of risk can vary a lot. Even the sign can change. There are many possible permutations depending on the window size used for measuring volatility and the difference in the timing of shocks.

Take the two cities in the above example that each receive positive shocks, one in period one and the other in period five. The city that receives the shock in period five is "riskier" (from the perspective of the econometrician using a five year window for volatility) for the years five through nine than the city receiving the shock in period one. A panel regression of homeownership rates on prices and volatility using years 5 through, say, 20 would find a positive coefficient on volatility (and in this case, the coefficient on prices would not be identified as there would be no cross-sectional variation in prices). In other words, a regression with a sample that contained a period of above-trend price growth for many cities, where the timing of this above trend growth differed

over cities, could easily have an upwardly biased coefficient on volatility. Likewise, a regression with a sample period with differentially timed, below-trend price growth for many cities would have downwardly biased coefficients on volatility.[43]

7 Conclusion

There is significant variation across cities in homeownership, LTV, house prices and housing risk. Much of the variation in house prices and housing risk has a common source - variation in the value or scarcity of land. This makes regression-based methods of separating the effects of risk from prices on household behavior difficult and inconclusive. Instead, we build a life-cycle model of homeownership which we match to the mean city in our data.

The model is able to explain much of the cross-city variation in homeownership and LTV and matches the variation in the data particularly well for younger households. We find that it is the relatively higher prices in cities with scarce land which causes their lower homeownership rates, while it is their relatively higher volatility that causes homeowners in these cities to borrow less. So, we do not find that more risk leads households to own more. Instead, more risk leads to higher reliance on non-housing savings. This result highlights the importance of including other means of imperfect insurance in asset allocation models with incomplete markets.

The main question explored in this paper concerns homeownership and insurance. Several of the facts developed in this paper - the relationships between land values, land scarcity and house prices, and between housing availability (the housing ladder) and homeownership - beg further examination and a full structural explanation. It would also be worth using a general equilibrium version of the model to examine the different local effects of aggregate (country-wide) shocks. Finally, it would be useful to explore whether land scarcity, via LTV, can help explain any recent geographical patterns in mortgage defaults.

References

Albouy, D., Aug 2009a. The unequal geographic burden of federal taxation. Journal of Political Economy 117 (4), 635–667.

Albouy, D., May 2009b. What are cities worth? land rents, local productivity, and the capitalization of amenity values. NBER Working Papers 14981, National Bureau of Economic Research, Inc.

---

[43]For example, in the period 2006-2010, there is a general fall in house prices in the U.S. with cities like Las Vegas experiencing price declines sooner than cities like Seattle.

Attanasio, O., Banks, J., Meghir, C., Weber, G., 1999. Humps and bumps in lifetime consumption. Journal of Business and Economics.

Attanasio, O. P., Weber, G., 1989. Consumption and saving: Models of intertemporal allocation and their implications for public policy. Journal of Economic Literature 48 (3), 693–751.

Banks, J., Blundell, R., Oldfield, Z., Smith, J. P., September 2010. House price volatility and the housing ladder. Tech. rep., Institute For Fiscal Studies.

Ben-Shahar, D., March 1998. On the optimality of the hybrid tenure mode. Journal of Housing Economics 7 (1), 69–92.

Blundell, R., Pistaferri, L., Preston, I., December 2008. Consumption inequality and partial insurance. American Economic Review 98 (5), 1887–1921.

Browning, M., Lusardi, A., 1996. Household saving: Micro theories and micro facts. Journal of Economic Literature 34 (4), 1797–1855.

Cagetti, M., 2003. Wealth accumulation over the life cycle and precautionary savings. Journal of Business & Economic Statistics 21 (3), 339–353.

Campbell, J., Cocco, J., November 2003. Household risk management and optimal mortgage choice. Quarterly Journal of Economics 118 (4), 1449–1494.

Campbell, S. D., Davis, M. A., Gallin, J., Martin, R. F., September 2009. What moves housing markets: A variance decomposition of the rent-price ratio. Journal of Urban Economics 66 (2), 90–102.

Caplin, A., Chan, S., Freeman, C., Tracy, J., 1997. Housing Partnerships. MIT Press.

Chaney, T., Sraer, D., Thesmar, D., 2012. The collateral channel: How real estate shocks affect corporate investment. American Economic Review Forthcoming.

Chiuri, M. C., Jappelli, T., 2003. Financial market imperfections and home ownership: A comparative study. European Economic Review 47 (5), 857 – 875.

Cocco, J., 2005. Portfolio choice in the presence of housing. Review of Financial Studies 18 (2).

Coulson, N. E., Fisher, L. M., 2009. Housing tenure and labor market impacts: The search goes on. Journal of Urban Economics 65 (3), 252 – 264.

Davidoff, T., March 2006. Labor income, housing prices and homeownership. Journal of Urban Economics 59 (2), 209–235.

Davis, M. A., Ortalo-Magne, F., April 2011. Household expenditures, wages, rents. Review of Economic Dynamics 14 (2), 248–261.

Davis, M. A., Palumbo, M. G., January 2008. The price of residential land in large US cities. Journal of Urban Economics 63 (1), 352–384.

Diaz, A., Luengo-Prado, M. J., 2008. On the user cost and homeownership. Review of Economic Dynamics 11 (3), 584–613.

Favilukis, J., Ludvigson, S. C., Van Nieuwerburgh, S., May 2010. The macroeconomic effects of housing wealth, housing finance, and limited risk-sharing in general equilibrium. NBER Working Papers 15988, National Bureau of Economic Research, Inc.

Fernandez-Villaverde, J., Krueger, D., June 2007. Consumption over the life cycle: Facts from consumer expenditure survey data. The Review of Economics and Statistics 89 (3), 552–565.

Fisher, J. D. M., Gervais, M., Aug 2011. Why has home ownership fallen among the young? International Economic Review 52 (3), 883–912.

Flavin, M., Yamashita, T., March 2002. Owner-occupied housing and the composition of the household portfolio over the life cycle. American Economic Review 92 (1).

Genesove, D., Mayer, C. J., June 1997. Equity and time to sale in the real estate market. American Economic Review 87 (3), 255–69.

Gervais, M., 2002. Housing taxation and capital accumulation. Journal of Monetary Economics 49, 1461–1489.

Gourinchas, P.-O., Parker, J. A., January 2002. Consumption over the life cycle. Econometrica 70 (1), 47–89.

Halket, J., Vasudev, S., Dec. 2011. Home ownership, mobility and savings over the life cycle.

Han, L., September 2008. Hedging house price risk in the presence of lumpy transaction costs. Journal of Urban Economics 64 (2), 270–287.

Han, L., November 2010. The effects of price risk on housing demand: Empirical evidence from U.S. markets. Review of Financial Studies 23 (11), 3889–3928.

Hryshko, D., Luengo-Prado, M. J., Sorensen, B. E., Jun. 2010. House prices and risk sharing. CEPR Discussion Papers 7894, C.E.P.R. Discussion Papers.

Hurst, E., Stafford, F., December 2004. Home is where the equity is: Mortgage refinancing and household consumption. Journal of Money, Credit and Banking 36 (6), 985–1014.

Jeske, K., Krueger, D., 2005. Housing and the macroeconomy: the role of implicit guarantees for government-sponsored enterprises. Working Paper 2005-15, Federal Reserve Bank of Atlanta.

Kreps, D. M., Porteus, E. L., January 1978. Temporal resolution of uncertainty and dynamic choice theory. Econometrica 46 (1), 185–200.

Lam, K., Kaul, B., April 2003. Analysis of housing finance issues using the american housing survey (ahs). Tech. rep., U.S. Department of Housing and Urban Development, PD&R.
URL `http://www.huduser.org/publications/pdf/AhsAnalysis.pdf`

Lamont, O., Stein, J., 1999. Leverage and house-price dynamics in us cities. RAND Journal of Economics 30, 498–514.

Lazear, E. P., Michael, R. T., 1980. Family size and the distribution of real per capita income. American Economic Review 70, 91–107.

Leth-Petersen, S., 2010. Intertemporal consumption and credit constraints: Does total expenditure respond to an exogenous shock to credit? American Economic Review 100 (3), 1080–1103.

Li, W., Yao, R., August 2007. The life-cycle effects of house price changes. Journal of Money, Credit and Banking 39 (6), 1375–1409.

Lustig, H., Van Nieuwerburgh, S., April 2010. How much does household collateral constrain regional risk sharing? Review of Economic Dynamics 13 (2), 265–294.

Martin, R. F., 2003. Consumption, durable goods, and transaction costs. International Finance Discussion Papers 756, Board of Governors of the Federal Reserve System (U.S.).

Mian, A., Sufi, A., August 2011. House prices, home equity-based borrowing, and the US household leverage crisis. American Economic Review 101 (5), 2132–56.

Nordvik, V., December 2001. A housing career perspective on risk. Journal of Housing Economics 10 (4), 456–471.

Orshansky, M., 1965. Counting the poor: Another look at the poverty profile. Social Security Bulletin January, 3–29.

Ortalo-Magne, F., Rady, S., 2002. Tenure choice and the riskiness of non-housing consumption. Journal of Housing Economics 11, 266–279.

Ortalo-Magne, F., Rady, S., 2006. Housing market dynamics: On the contribution of income shocks and credit constraints. Review of Economic Studies 73, 459–485.

Paciorek, A., Sinai, T., 2012. Does home owning smooth the variability of future housing consumption? Journal of Urban Economics 71 (2), 244 – 257.

Paciorek, A. D., 2012. Supply constraints and housing market dynamics. Finance and Economics Discussion Series 2012-01, Federal Reserve Board.

Piazzesi, M., Schneider, M., Tuzel, S., 2007. Housing, consumption and asset pricing. Journal of Financial Economics 83, 531–569.

Piketty, T., Saez, E., 2007. How progressive is the US federal tax system? a historical and international perspective. Journal of Economic Perspectives 21, 3–24.

Rosen, H. S., Rosen, K. T., Holtz-Eakin, D., 1984. Housing tenure, uncertainty and taxation. Review of Economics and Statistics 66, 405–416.

Ruggles, S., Alexander, J. T., Genadek, K., Goeken, R., Schroeder, M. B., Sobek, M., 2010. Integrated public use microdata series: Version 5.0 [machine-readable database]. Minneapolis: University of Minnesota.

Saiz, A., 2010. The geographic determinants of housing supply. Quarterly Journal of Economics 125 (3), 1253–1296.

Sinai, T., Souleles, N. S., May 2005. Owner-occupied housing as a hedge against rent risk. The Quarterly Journal of Economics 120 (2), 763–789.

Stein, J., 1995. Prices and trading volume in the housing market: A model with down-payment effects. Quarterly Journal of Economics 110, 379–406.

Storesletten, K., Telmer, C., Yaron, A., 2004a. Consumption and risk sharing over the life cycle. Journal of Monetary Economics 51 (3), 609–633.

Storesletten, K., Telmer, C. I., Yaron, A., June 2004b. Cyclical dynamics in idiosyncratic labor market risk. Journal of Political Economy 112 (3), 695–717.

Tauchen, G., 1986. Finite state markov-chain approximations to univariate and vector autoregressions. Economics Letters 20, 177–181.

US Department of Commerce, B. o. t. C., 1991. Ser. p-60, no. 176-rd, table f and app. a. In: Current Population Reports. US Government Printing Office.

Verbrugge, R., December 2008. The puzzling divergence of rents and user costs, 1980-2004. Review of Income and Wealth 54 (4), 671–699.

Vissing-Jorgensen, A., Attanasio, O. P., 2003. Stock-market participation, intertemporal substitution, and risk-aversion. American Economic Review 93 (2), 383–391.

Yao, R., Zhang, H. H., 2005. Optimal consumption and portfolio choices with risky housing and borrowing constraints. Review of Financial Studies 18 (1).

Yogo, M., 04 2006. A consumption-based explanation of expected stock returns. Journal of Finance 61 (2), 539–580.

A   Robustness to Supply-side effects

A.1   Background

First, we briefly outline the different types of mortgage available in the United States [44]. Mortgage loans may either be conventional or non-conventional. To qualify for a conventional loan, households must pass credit and PITI (income) tests. If they cannot afford a threshold down payment (often as high as 20%), they will also have to purchase PMI (private mortgage insurance) to qualify for a conventional loan. Non-conventional loans are guaranteed by the government, through the Federal Housing Administration (FHA) or Department of Veteran Affairs (VA). They tend to be more appropriate for households who require a large LTV.

---

[44]For more detail see, e.g., Caplin et al. (1997).

Conventional loans may either be conforming or nonconforming. Loans are conforming if they fall below a dollar threshold, which varies with time. Until 2008, this threshold was nationally uniform (our sample excludes years after 2008). Conforming loans are subject to cheaper rates, because they are more liquid: Fannie Mae or Freddie Mac will provide guarantees enabling a lender to sell them to the secondary market.

A.2 Conforming and non-conforming loans

Here, we consider the impact of conforming loan limits. The existence of this nationally uniform loan limit may well be responsible for our LTV result. In more expensive cities, the conforming loan limit is more likely to bind. As a result, households will be forced to make a larger down payment (to qualify for the cheaper rates on conforming loans). And, this will yield a negative correlation between price and LTV (and consequently, between price risk and LTV too). If the conforming loan limit is driving our results, then the effect should be stronger the closer the loan size is to the conforming loan limit: increases in land value would be less likely to lead to increases in loan value if that means the household will go over the conforming loan size limit.

To test whether the loan limit is driving this effect, we check the LTV-price level/volatility correlation in samples delineated by the ratio of loan size to loan limit (restricting our attention to households with conventional mortgages). It turns out, though, that the correlation is strongly negative (especially) for loans well below the limit and less so for loans close to the limit. So, we conclude that the loan limit cannot be responsible for the correlation.

The results are reported in Table 7. In each case, observations are at household-level, and the dependent variable is LTV. The regressor of interest is price volatility in Panel A and log house price in Panel B. Also included are a range of household-level controls (see table notes) and a full set of purchase year fixed effects (we only include households that purchased their home between 1990 and 2004).

Regressions are disaggregated into samples delineated by loan-to-limit ratio. Also, the sample is restricted to households with conventional mortgages (we are interested in the impact of the conforming loan limit). The loan-to-limit ratio for each sample is reported at the top of the columns (0-0.25, 0.25-0.5, 0.5-0.75, 0.75-1, 1-1.25, 1.25-1.5 and >1.5).

The effect on LTV is negative in all samples for both price volatility and levels. For volatility (Panel A), the effect is very large and statistically significant for the 0-0.25 sample (-4.9) and the 0.25-0.5 sample (-1.6). But, the effects for all the other samples fall below 0.8 and are statistically insignificant. The effects around the conforming loan limit (0.75-1 and 1-1.25) actually tend to be *smaller* than elsewhere. This suggests that the negative effect is not being driven by some

interaction with the conforming loan limit. The patterns are very similar for price level in Panel B, though more of the samples are statistically significant.

A.3 Local variation in effective interest rates

An alternative hypothesis is that banks subject households in riskier cities to higher mortgage interest rates - and this could explain the lower ownership rate in these cities. Similarly, it could explain why households in these cities choose to take out smaller loans (relative to home value). Interest rates may vary across cities because of differences in state-level regulation.

However, it turns out that effective interest rates[45] are actually *lower* in expensive/risky cities. This is illustrated by Figure 16, using data from the MIRS on 25 major cities. There is a strong negative correlation across cities between the interest rate and both price risk and level (R squared is 40-50% in each case). The interest rate varies from 7% in the most expensive/risky cities to 8% in the least. And so, this cannot explain why ownership rates and LTV ratios are also lower in expensive cities.

B Parametrization

B.1 Household life-cycle and preferences

We calibrate the discount factor, $\beta = 0.95$, housing's share in the utility function, $\sigma = 0.3$ following Favilukis et al. (2010) and the inverse of the intertemporal elasticity of substitution, $\phi = 5$, following Piazzesi et al. (2007). Estimates of risk aversion vary widely, particularly when the parameter is separately identified from the intertemporal elasticity of substitution. Some studies have point estimates with $\gamma = 20$ or higher but with equally large confidence intervals (see Attanasio and Weber (1989); Vissing-Jorgensen and Attanasio (2003), and, for values over 100, Yogo (2006)). Since such a large value of γ would imply an outlandish level of precautionary savings in our model, we choose $\gamma = 3$, which is well within the more traditional range of two to five that most studies prefer (see Lustig and Van Nieuwerburgh (2010); Hryshko et al. (2010); Li and Yao (2007)).

---

[45]"Effective" because it accounts for any up-front fees

B.1.1 Family size equivalence

We collect data from the period 1970-1993 in the CPS. We control for year effects by using year dummies. The family size profile is generated by the following regression:

$$F_{iat} = \sum_{k=21}^{81} \beta_k 1_k + \sum_{t'=1970}^{1993} \beta_{t'} 1_{t'} + \varepsilon_{iat}$$

where 1_k is a year dummy which takes on value 1 when $a = k$, and $1_{t'}$ is the year dummy that takes on value 1 when $t' = t$.

Figure 17 shows the profiles of family size from the CPS. Family size increases sharply when the household is young, peaking at age 39.

In order to adjust the household's housing and consumption stream, we use a household equivalence scale. The objective of an equivalence scale is to measure the change in consumption needed to keep the welfare of the family constant as the family size varies. Note that using per capita consumption assumes that the family converts consumption expenditure into utility flow following constant returns to scale. Lazear and Michael (1980) point to the existence of family goods, economies of scale and complementarities, which are all factors that they show to be significant. We therefore use a household equivalence scale that is not constant returns to scale. Table B.4 lists some equivalence scales. L-M stands for Lazear and Michael (1980), US Dept of Commerce refers to US Department of Commerce (1991) and F-V&K stands for Fernandez-Villaverde and Krueger (2007). Lazear and Michael's scale takes greater account of common or public goods, so that the impact of family size is less then other equivalence scales (compare, for instance, Orshansky (1965)). We use the housing equivalence scale used by Fernandez-Villaverde and Krueger (2007).

All households in the model economy have the same life-cycle profile of family size, which is set to the average family size at each age in the CPS. To account for non-integer family sizes, we assume that the adjustment factor is linear within the family sizes specified in Table (B.4). Figure 17 shows the equivalent, normalized family size over the life cycle.

B.2 Assets

We set the down payment requirement, $d = .1$. We set the transaction cost of buying, $\theta_b = .08$, within the range typically chosen by the literature (Martin (2003); Fisher and Gervais (2011)). We set the interest rates $r_l = .04$. The average difference between a 30 year fixed rate mortgage and the 30 year U.S. Treasury bond is between 1 percent and 2 percent for 1977-2010 so we set $r_b = .06$.

B.2.1 Initial wealth distribution

We calibrate the wealth distribution of newborns using the distribution of wealth among 21-25 year olds in the Survey of Consumer Finances (SCF) waves from 1989-2001. We drop top-coded observations and households with negative wealth and students from the sample and use the sample weights provided by the SCF. We parametrize the initial wealth distribution as an exponential distribution. That gives us one parameter that we have to match.

$$f(b_0) = \lambda_w e^{-\lambda_w b_0}$$

where b_0 is the initial wealth, and λ_w is the parameter to estimate in the exponential distribution. We estimate λ_w by matching the mean of the initial wealth distribution.

$$\lambda_w = \frac{1}{\overline{b_0}}$$

This gives us $\lambda_w = 0.00589$. We convert the initial wealth distribution in the data to model terms by scaling by the ratio of average labor earnings at age 21 in the model to average labor earnings at age 21 in the data.

B.3 Taxes

There are two forms of taxes in the model economy - income tax, t_y, and property tax, t_p. Piketty and Saez (2007) uses public use micro-files of tax return data from the Internal Revenue Service, which have the advantage of being aggregated to the household level already. The income tax rate we choose, $t_y = 0.2$, is in the same range that they compute for the US economy[46].

   We use data from the IPUMS 1990 5 percent sample. The variables used are the amount of property tax paid and the estimated value of the house. We remove top-coded variables from the sample, and consider only owner-occupiers. Sample observations are weighted using the household weights given in the data set. The weighted average of the ratio of the amount of property tax paid to the estimated value of the house is 0.012. In the model we set $t_p = 0.01$.

B.4 Earnings Process

We parametrize the idiosyncratic and age-profile portion of the household's earnings following Halket and Vasudev (2011), who estimate a process similar to Storesletten et al. (2004a) but also

---

[46]See Table 1, page 6 in their paper

control for regional variability (in their case, at the U.S. state level) in earnings rather than just national variability. We set the standard deviation of idiosyncratic innovations, $\sigma_\psi = .098$ and let the initial (fixed effect) distribution have a standard deviation of 0.5 (since the persistent component of earning follows a random walk, a fixed effect is equivalent to households entering at age 21 with a value ψ_{21}^i drawn from normal distribution with standard deviation 0.5). As is well known, the variance of the transitory shock is not separately easily identified from the variance of measurement error in these approaches to estimation. We set $\sigma_\rho = .25$, which is within bounds found by Storesletten et al. (2004b); Blundell et al. (2008) We discretize the innovations with a 3-point distribution following Tauchen (1986).

We set the pension at 60 percent of final earnings, $\zeta = 0.6$.

Figures and tables
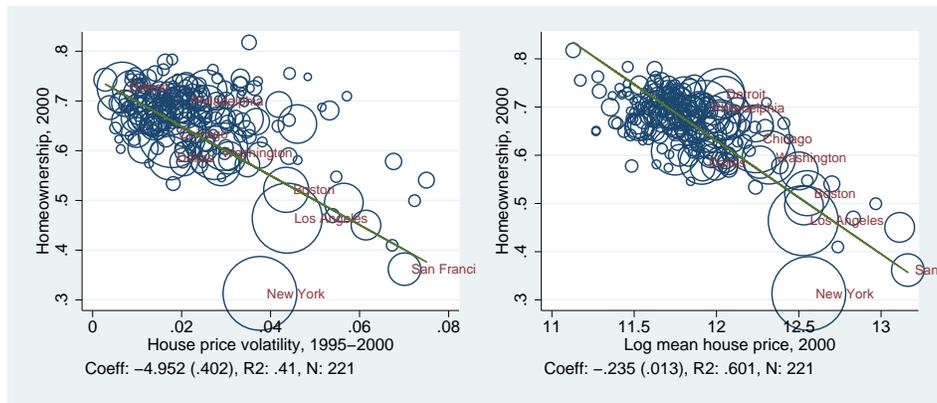
Figure 1: Homeownership, price risk and price levels
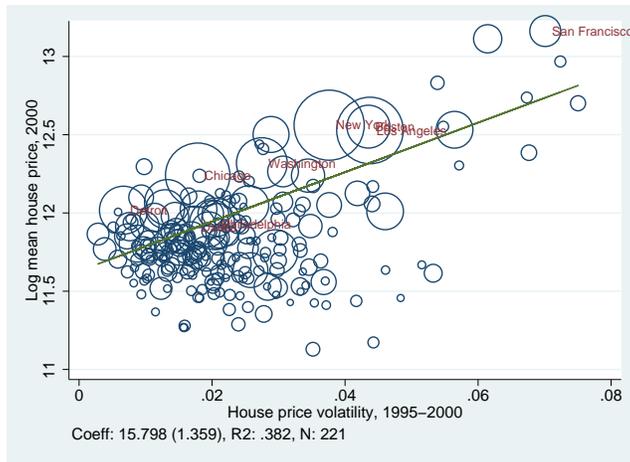
Figure 2: Price levels and risk
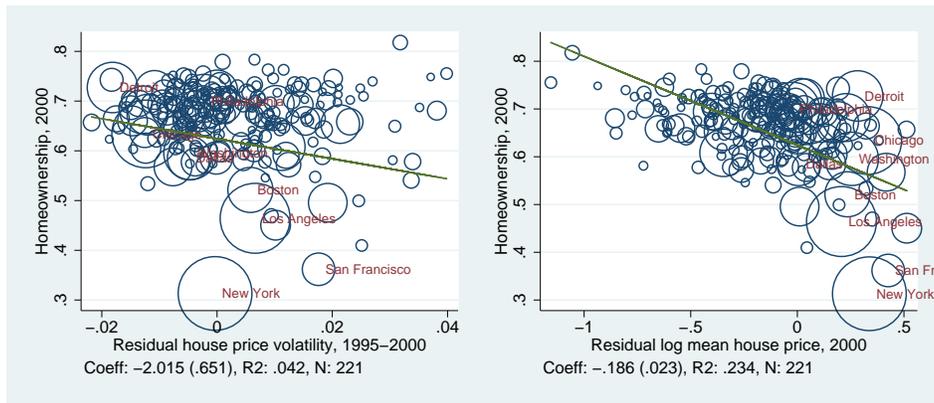
Figure 3: Homeownership and residual price levels and volatilities
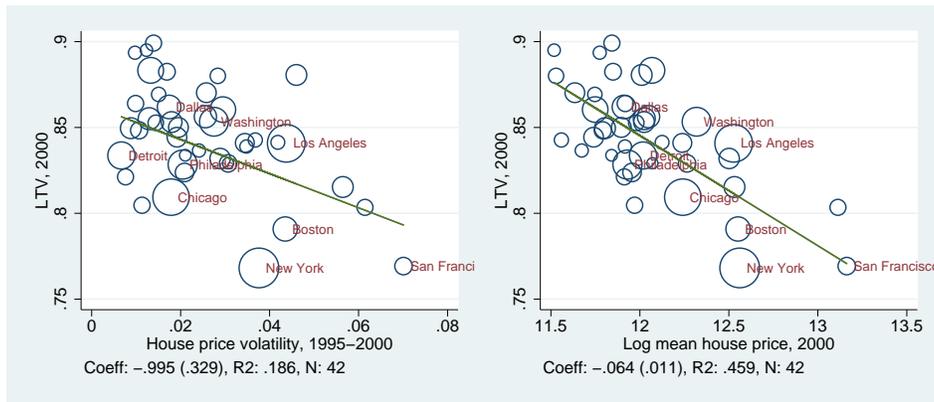
Figure 4: LTV (AHS), price risk and price levels
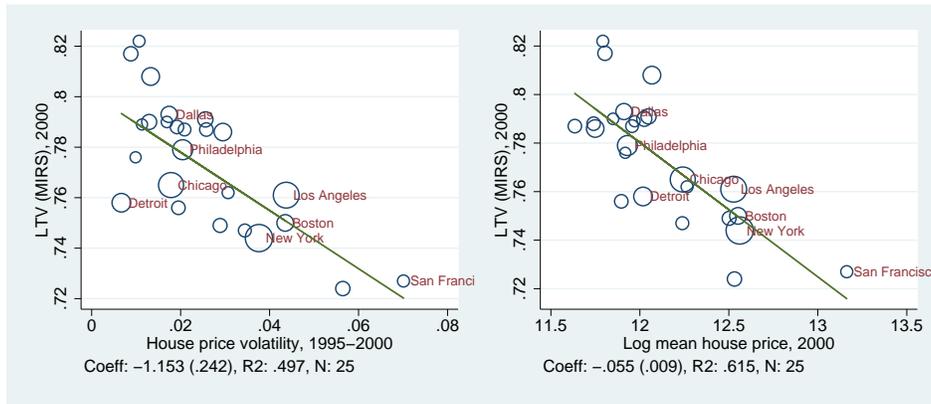
Figure 5: LTV (MIRS), price risk and price levels
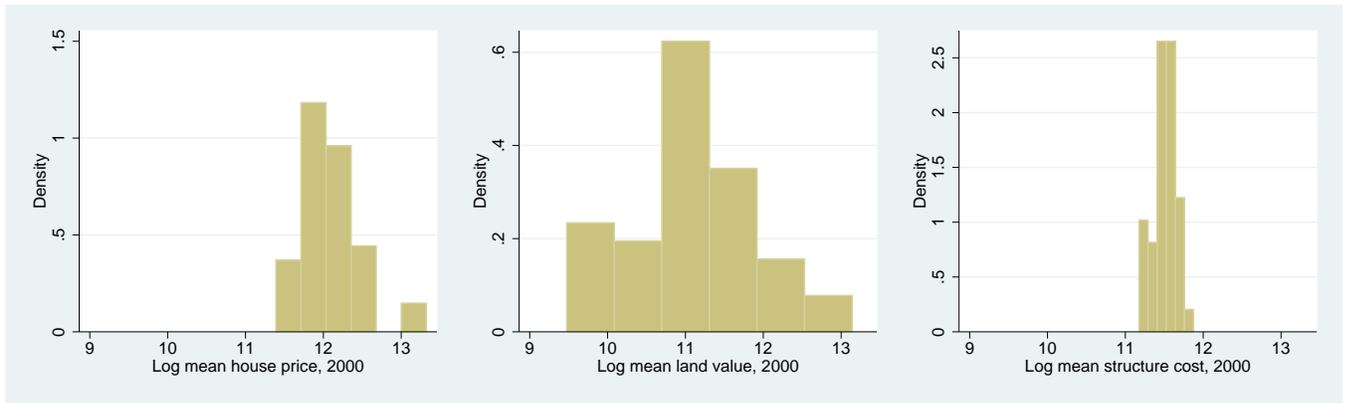

Figure 6: Histograms of city price, structure cost and land value levels
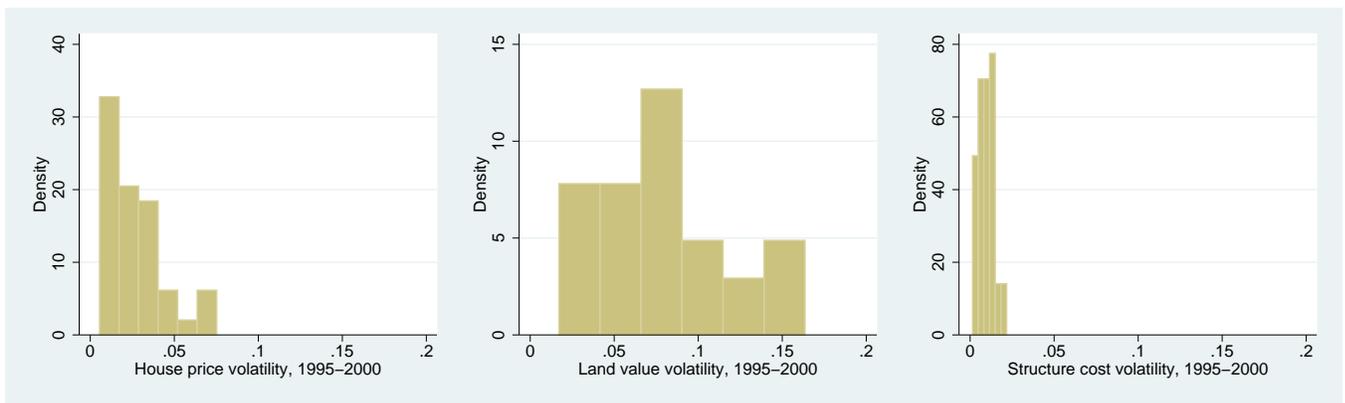
Figure 7: Histograms of city price, structure cost and land value volatilities
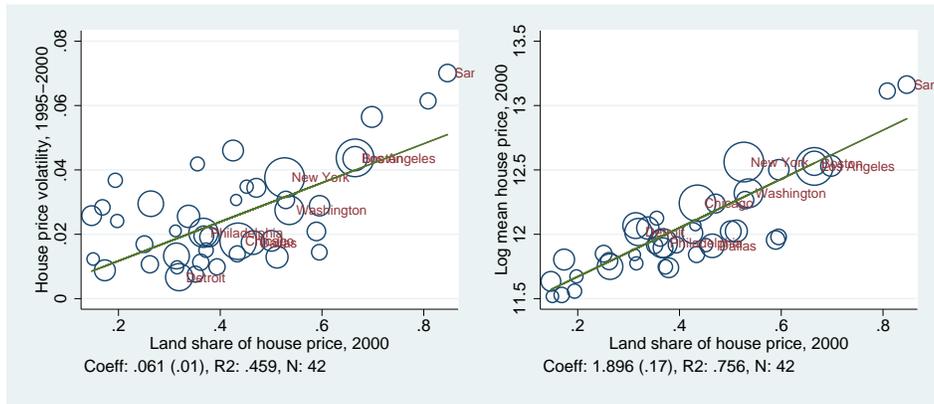
Figure 8: Price risk and levels and land share
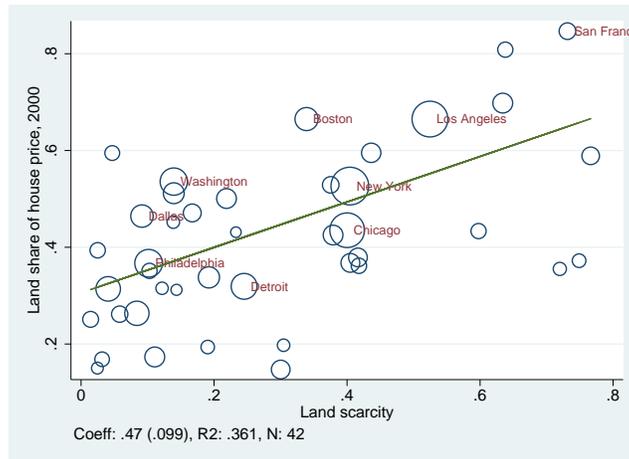

Figure 9: Land share and land scarcity
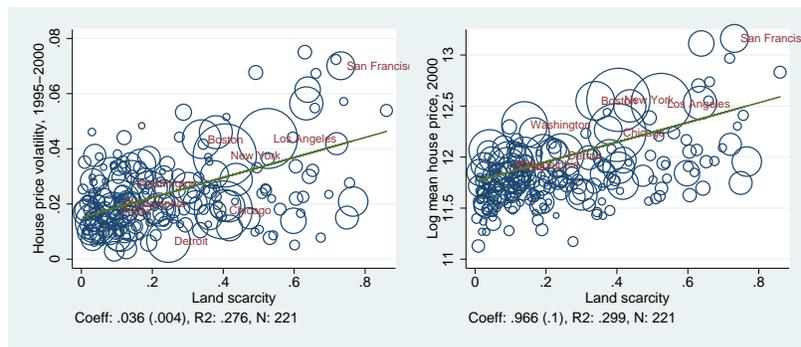
Figure 10: Price risk and levels and land scarcity

41
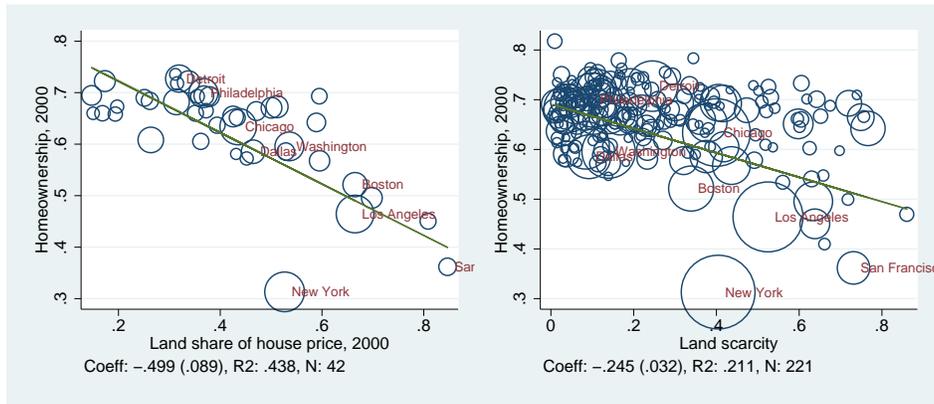
Figure 11: Homeownership, land share and land scarcity
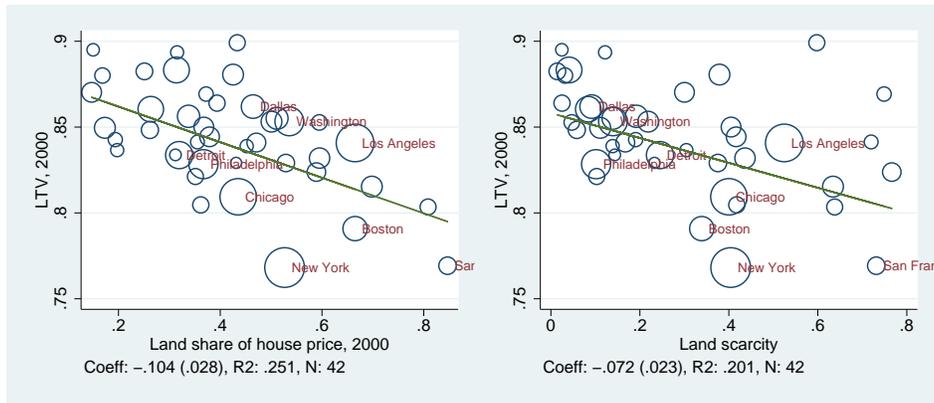
Figure 12: LTV, land share and land scarcity
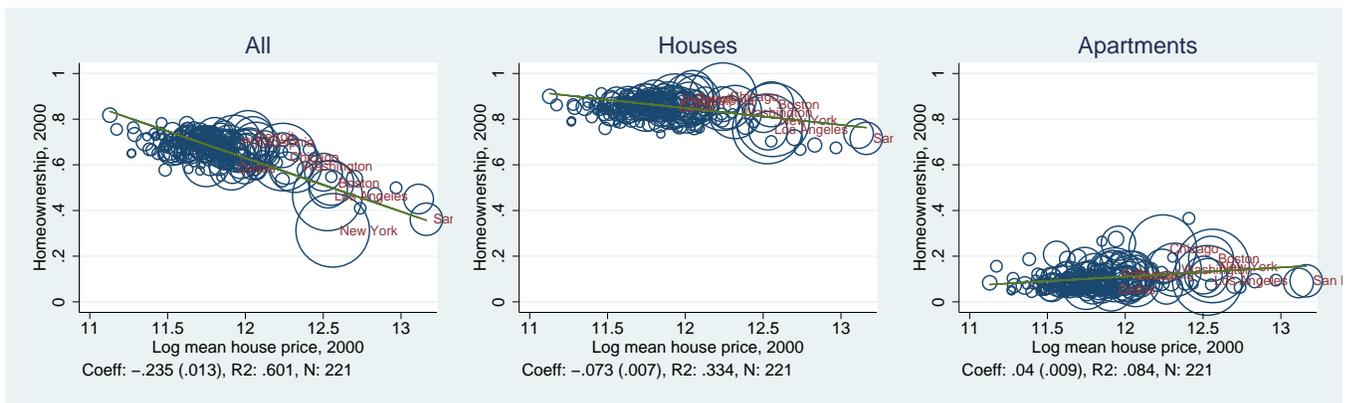
Figure 13: Homeownership-price relationship by dwelling type
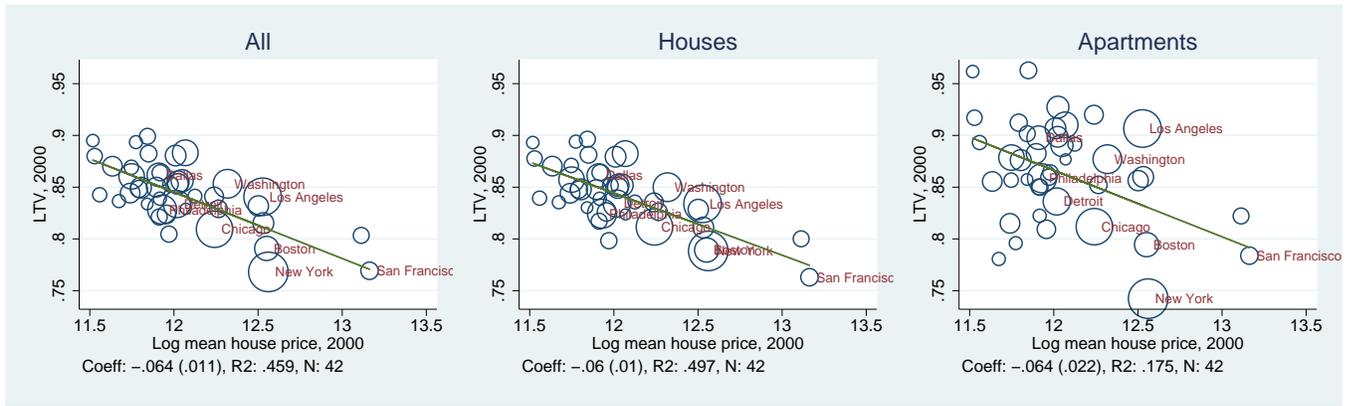
Figure 14: LTV-price relationship by dwelling type
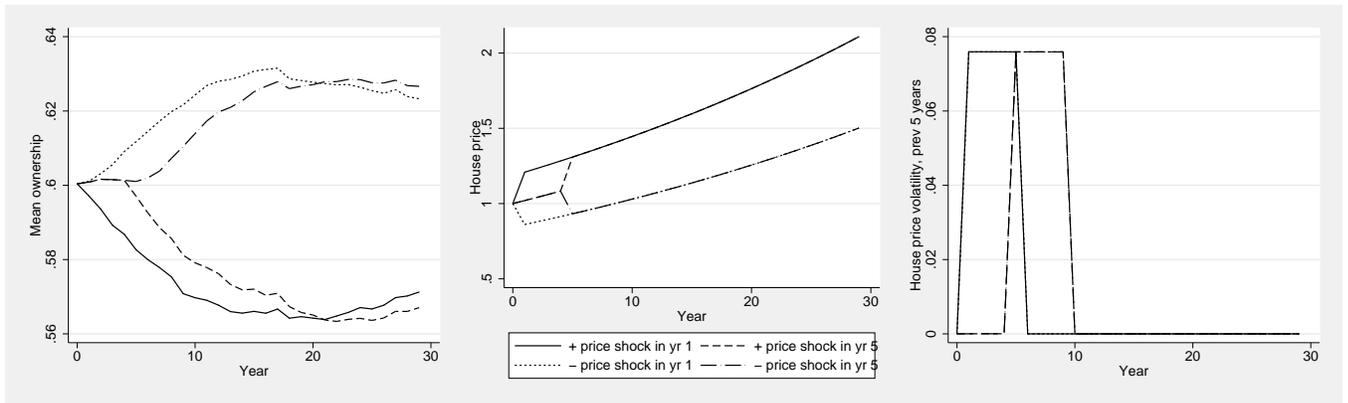
Figure 15: Impulse responses: mean ownership, house price level and volatility
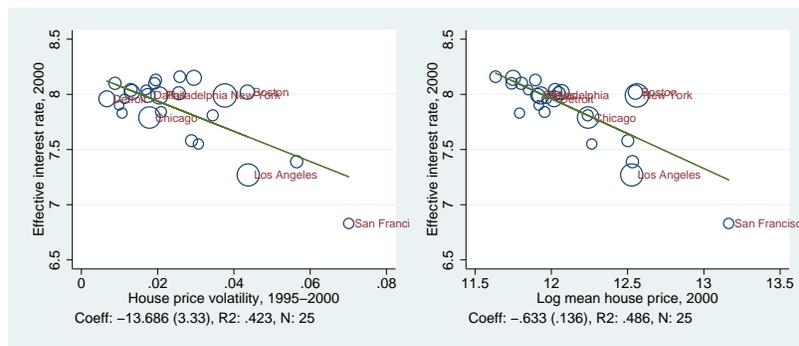
Figure 16: Effective mortgage interest rate, price risk and price levels
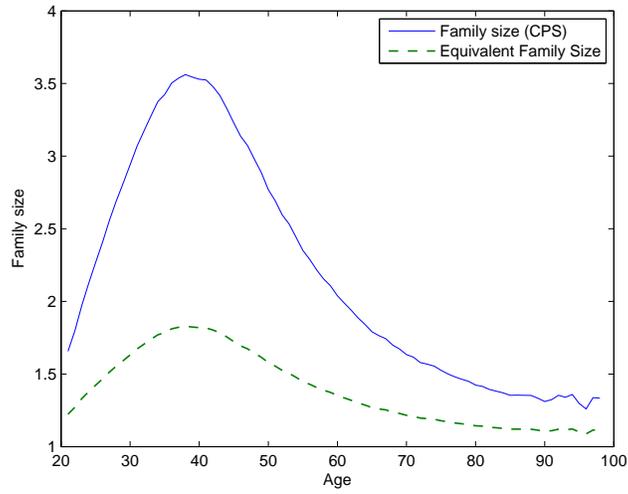
Figure 17: L-M adjusted family size profile (yearly bins)

Table 1: Invariant parameters

Parameter	Description	Value
β	Discount factor	0.95
σ	Housing's share in utility	0.30
ϕ	IIES	5
γ	Risk Aversion	3
t_y	Income tax	0.20
t_p	Property tax	0.01
σ_ρ	Std. dev. of the idiosyncratic transitory shock	0.25
σ_ψ	Std. dev. of the idiosyncratic permanent shock	0.098
r_b	Interest rate on loans	6%
r_l	Interest rate on savings	4%
d	Down payment	0.1
θ_b	Home buyer's transaction cost	0.08
ζ	Replacement rate for pensions	0.6

Table 2: Matched parameters

Parameter	Description	Mean value	Interquartile range
σ_v	Std. dev of shock to wages (corr with house prices)	0.012	0.010-0.013
σ_ε	Std. dev of idiosyncratic shock to house prices	0.022	0.015-0.025
$\underline{h}^o$	Min owner-occupied house size	4	4 - 4
$\overline{h}^r$	Max rental house size	8.25	8.25 - 8.25
δ_j	Housing maintenance	0.0112	0.012-0.010
p_{2000}	Price level in 2000	1	0.81 - 1.09
W_{2000}	Wage level in 2000	1	0.96 - 1.02

Price and wage levels are normalized so that prices and wages are equal to one for all cities with the average level of land scarcity in the year 2000. The units on the house size parameters are median household earnings for 21 year olds, and these parameters are not changed across cities.

Table 3: Model fit

	Data source	City sample	Data: 221 cities	Model
Homeownership rate	Census	221	0.62	0.63
Homeownership rate under 65	Census	221	0.61	0.61
Owned/rented home size ratio	AHS	42	2.07	2.05
LTV	AHS	42	0.84	
LTV (conventional loans only)	AHS	42	0.80	0.71
LTV (conventional loans only)	MIRS	25	0.77	

This table compares key parameters in the data with the model. The first column shows the number of cities on which the data estimates are based (see Section 2.1 for further details). The second column gives the mean (weighted by census sample size) for the relevant variable across those cities (NB restricting the larger samples to 25 cities has only a negligible effect on the estimated means). The AHS LTV estimates are conditional on household characteristics (this is not trivial because of the omitted top-coded observations); see Section 2.1 for further details and estimation procedure. We also report the mean LTV across cities for the sample of conventional loans in the AHS. This makes it more comparable with the LTV estimate from the MIRS (the final row), whose sample excludes non-conventional loans (see Section 2.1).

Table 4: Homeownership profile: data and model

Age	Data source	City sample	Data: mean	Model: mean
21-35	Census	221	0.38	0.25
36-50	Census	221	0.67	0.72
51-65	Census	221	0.76	0.88
66-75	Census	221	0.78	0.68

See notes under Table 3. This table reports mean ownership rates by age group.

Table 5: Slopes with respect to land scarcity

age	Own: Data	Own: Model	LTV: Data (AHS)	LTV: Data (MIRS)	LTV: Model
all ages	-0.25**	-0.20†	-0.07**	-0.08**	-0.06†
21-35	-0.23**	-0.32	-0.09**	N/A	-0.16
36-50	-0.24**	-0.22†	-0.05**	N/A	0.04
51-65	-0.18**	-0.09	-0.08	N/A	0.22
66-75	-0.17**	-0.13†	0.02	N/A	0.00†

This table compares cross-city slopes of ownership rates and LTV with respect to land scarcity, for both the data and model. For the data, reported coefficients are taken from cross-city OLS regressions (weighted by census sample size) of mean ownership or LTV (for the age group in question) on land scarcity. The local ownership and AHS LTV estimates are conditional on observed household characteristics (see Section 2.1 for estimation procedure), but not the MIRS. Also, there is no available disaggregation of the MIRS data by age group. ** signifies that the estimate from the data is significant at the 95 percent confidence level. † signifies that the model estimate falls within the 95 percent confidence interval of the data.

Table 6: Contribution from various elements: slopes with respect to land scarcity

Parameter	Homeownership: slopes	LTV: slopes
σ_v	-0.02	-0.01
σ_ε	-0.03	-0.01
δ_j	-0.01	0.01
p_{2000}	-0.23	0.02
W_{2000}	0.04	0.00

Table 7: Regressions of LTV on price volatilities and levels, for samples delineated by loan-limit ratio

PANEL A: PRICE VOLATILITIES

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Sample: loan/limit	0-0.25	0.25-0.5	0.5-0.75	0.75-1	1-1.25	1.25-1.5	>1.5
House price volatility	-4.882***	-1.593***	-0.587	-0.223	-0.552	-0.767	-0.474
	(1.653)	(0.586)	(0.392)	(0.483)	(0.380)	(0.455)	(0.863)
Household controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Purchase year effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,360	4,566	3,285	1,506	550	275	304
MSAs	42	42	42	42	42	42	42
R-squared	0.051	0.020	0.007	0.008	0.055	0.099	0.169

PANEL B: PRICE LEVELS

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Sample: loan/limit	0-0.25	0.25-0.5	0.5-0.75	0.75-1	1-1.25	1.25-1.5	>1.5
Log house price	-0.388***	-0.116***	-0.042***	-0.005	-0.031*	-0.042**	-0.055
	(0.051)	(0.017)	(0.014)	(0.021)	(0.017)	(0.019)	(0.033)
Household controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Purchase year effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,360	4,566	3,285	1,506	550	275	304
MSAs	42	42	42	42	42	42	42
R-squared	0.119	0.037	0.012	0.008	0.060	0.108	0.189

Regressions are run separately for samples delineated by loan-to-limit ratio (the ratio of loan size to conforming loan limit, reported above each column). For each sample, we separately estimate the effect of price volatility (in Panel A) and log price level (in Panel B) on household-level LTV. All regressions control for the household characteristics listed in Section 2.1 (i.e. those used to condition the local LTV estimates), as well as a full set of purchase year effects. We use the composite sample described in Section 2.1, though we exclude all household with non-conventional mortgages. There are 42 MSAs in the sample. SEs, clustered by city, in parentheses. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table 8: Family size equivalence scale

Family Size	L-M	Orshansky (1965)	US Dept of Commerce	F-V&K
1	100	100	100	100
2	106	126	128	134
3	128	151	157	165
4	147	189	201	197
5	169	223	237	227