

NBER WORKING PAPER SERIES

HOUSE PRICES AND RENTS IN THE 21ST CENTURY

Lara Loewenstein
Paul S. Willen

Working Paper 31013
<http://www.nber.org/papers/w31013>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
March 2023

Thanks to Gadi Barlevy, Daniel Cooper, Chris Foote, Peter Ganong, Ed Glaeser, Jaromir Nosal, Joe Nichols, and audiences at the Federal Reserve Mortgage Forum, Homer Hoyt, the Notre Dame Housing Conference, the Cleveland Fed, the 2019 St. Louis Fed Rental Conference, the 2019 AREUEA national conference, Ohio State University, the 2019 housing conference hosted by the Federal Reserve Bank of Atlanta Center for Housing and Policy and the Georgia State University Department of Real Estate, and the 2021 Midwest Economics Association Annual Conference for helpful comments and suggestions. The views in this paper are not necessarily those of the Federal Reserve Bank of Boston, the Federal Reserve Bank of Cleveland, or the Federal Reserve System. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2023 by Lara Loewenstein and Paul S. Willen. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

House Prices and Rents in the 21st Century
Lara Loewenstein and Paul S. Willen
NBER Working Paper No. 31013
March 2023
JEL No. E30,E32,H31,R30,R31

ABSTRACT

We study the joint evolution of prices and rents of residential property. After constructing rent and price indices for renter- and owner-occupied properties, we decompose the change in the price of occupant-owned property into (1) changes in rent, (2) changes in the relative prices of investor- and occupant-owned properties, and (3) changes in the price-rent ratio. Via a simple model, we link our decomposition to different sources of variation in house prices. We argue that while the 2000s boom was plausibly driven by exuberant expectations, the boom of the 2020s more likely resulted from a preference shock.

Lara Loewenstein
Federal Reserve Bank of Cleveland
Research Department
1455 E 6th St
Cleveland, OH 44114
lara.loewenstein@clev.frb.org

Paul S. Willen
Federal Reserve Bank of Boston
Research Department
600 Atlantic Avenue
Boston, MA 02210-2204
and NBER
paul.willen@bos.frb.org

1 INTRODUCTION

House prices were unusually volatile over the first two decades of the 21st century. As depicted in Figure 1a, real house prices rose 60 percent from 2000 to 2005 then fell almost 60 percent before rebounding to end more than 100 percent above their 2000 level. Contrast this with the last 25 years of the previous century, when prices never moved more than 20 percent from their 1975 base. For market participants, researchers, and policymakers, explaining these sharp price movements has proved difficult, both retrospectively and in real time.

The discussion at the June 2005 Federal Open Market Committee (FOMC) meeting illustrates the challenge of making sense of house price movements. Despite a 40 percent increase in house prices over the previous four years, the FOMC concluded in its published minutes that “[house] prices might be somewhat above the levels consistent with the underlying factors, but measuring the extent of any overvaluation either nationally or in regional markets posed considerable conceptual and statistical difficulties.” Transcripts show that the debate concerned comparisons of house prices and one “underlying factor”: rents. The statistical difficulties centered on whether prices of owner-occupied properties and rents on renter-occupied properties were comparable. The conceptual issue was that even if participants could agree that the price-rent ratio was historically high, they could not agree on whether it indicated a speculative asset-price bubble or was due to fundamental changes in the economy that had made owner-occupied property more attractive relative to renter-occupied property.

In this paper, we provide a framework for interpreting fluctuations in house prices using a new data set of residential-property transactions involving single-family and multifamily homes (homes housing two, three, or four families). The advantage of our data set is that it contains both renter- and owner-occupied properties, and it includes both rent and sale transactions. As a result, we can measure price growth on owner- and renter-occupied properties, and we can calculate a price-rent ratio using only renter-occupied properties.

These features enable us to conduct the following decomposition:

$$p_o = \underbrace{\frac{p_o}{p_r}}_{\text{Price-Price Ratio}} \times \underbrace{\frac{p_r}{\text{rent}}}_{\text{Price-Rent Ratio}} \times \text{rent}, \quad (1.1)$$

where p_o is the price of owner-occupied housing and p_r is the price of renter-occupied housing. The first term on the right-hand side, the “price-price” ratio, exploits the fact that we can identify sale transactions as renter- or owner-occupied. The second term is the price-rent ratio, measured only on properties identified as renter-occupied. That is, in computing the price-rent ratio, we use explicitly comparable properties.

The main empirical results from our decomposition appear in the bottom panel of Figure 1. They show that the sources of growth in house prices were highly variable. Early in the sample, the price-rent ratio and the price-price ratio were the determinants of growth, and there was relatively little rent growth. But in the latter part, rent growth became much more important, and in the boom of the 2020s, rent growth was the main driver.

To address the conceptual problems, we use a simple model of the housing market. Within the model's framework, we consider five shocks to the housing market.

First, all else being equal, a positive *interest rate shock* lowers the price-rent ratio. However, in our data, the relationship between interest rates and the price-rent ratio is unstable. From 2000 to 2004 and after 2015, interest rates fell, and the price-rent ratio rose. But from 2004 to 2015, the opposite occurred. We argue that the relationship between interest rates and the price-rent ratio is confounded by the correlation between interest rates and expectations about house price growth.

Second, according to theory, an increase in preferences for housing increases both rents and prices, leaving the price-rent ratio unchanged. Since most of the growth in 2021 involved rents and not the price-rent ratio, we conclude that a *preference shock* is the most plausible explanation.

Third, housing *supply shocks* also affect both rents and prices, not the price-rent ratio, and we argue that high supply growth prior to the global financial crisis (GFC) of 2008 and low supply growth in the aftermath explains the low then high rent growth depicted in the bottom panel of Figure 1.

Fourth, *expectation shocks* are changes in beliefs about price growth or, more generally, growth in the marginal utility of housing. Since they generate self-fulfilling price increases, expectation shocks are often what people have in mind when they talk about bubbles or exuberance. According to the theory, positive expectation shocks drive up the price-rent ratio. Accordingly, Figure 1 indicates that expectation shocks are a plausible explanation for the boom of the 2000s but not for the boom of the 2020s.

Fifth, *reallocation shocks* redistribute housing across the population of households and investors and change the identity or the holdings of the marginal investor or the marginal homeowner. The leading example is a financial innovation that allows marginal households to spend more on housing. We show that the connection between reallocation shocks and prices and rents is very sensitive to modeling assumptions. In standard models, a reallocation shock raises rents and the price-price ratio and has no effect on the price-rent ratio, making reallocation shocks an unlikely explanation for the 2000s boom. However, we then consider richer models, such as that of Geanakoplos (2010), in which reallocation can generate more realistic predictions.

We contribute to three strains of the literature. First, we add to the literature on the

measurement of the price-rent ratio. Many early studies use Census Bureau or Bureau of Labor Statistics (BLS) data to impute rent for owner-occupied homes.¹ More recent studies focus on creating price-rent ratios on very similar properties. For example, Smith and Smith (2006) use Multiple Listing Service (MLS) data for 10 markets, matching rental and sales listings on property characteristics, and Davis et al. (2008) impute rents for owner-occupied housing using census microdata. Pancak (2017) uses data from Zillow on aggregated price-rent ratios based on estimated market values and rents for the same properties within specific geographic areas. However, these methods are imperfect.² Bracke (2015) uses property-level price-rent ratios on a sample of properties located in London.

Second, we add to the literature on the origins of the 2000s housing boom, which includes, but is not limited to, Mian and Sufi (2009), Foote et al. (2021), Kaplan et al. (2020), Favilukis et al. (2017), Geanakoplos (2010), and Greenwald and Guren (2021). Specifically, we describe the results of this literature in the context of the five shocks discussed above. Finally, we add to the literature on price-rent ratios in the context of house price expectations and market efficiency, which includes Case and Shiller (1990), Mankiw and Weil (1989), and Campbell et al. (2009).

2 DATA

We use property-level data from CoreLogic. These data come from public records, the Multiple Listing Service (MLS), and local government tax assessment files.

The MLS data contain information on both sale and rental listings and are collected from participating regional boards of Realtors that contribute to a centralized database. More than 90 boards participate, providing coverage for approximately 56 percent of all active MLS listings nationwide. For many markets, we have data going back to at least the early 2000s, and in select markets, we have data going back to the mid-1990s. The data set contains all the information from the listing, including the property street address, physical characteristics of the property (square footage of living space, number of bedrooms, bathrooms, etc.), and both the list and closing prices and rents. We supplement the MLS data with public records data—which contain information on deeds reflecting legal transfers of property and includes sale transactions that are not listed in the MLS data—and tax assessment files, which include parcel-level information on property tax collections.

We limit our attention to listings and transactions that are closed and that provide a closing price or rent. We remove all non-arm’s-length transactions and any properties

¹See Carson et al. (2006) and Verbrugge and Poole (2010) for a detailed discussion of these data.

²Glaeser and Gyourko (2008) highlight several potential measurement flaws in conventional price-rent ratio data.

identified as mobile homes, student housing, military housing, cooperatives,³ or mixed use. We use the descriptive text field in the MLS data to flag properties that recently were renovated or depreciated significantly.⁴ We remove the few properties that sold for less than \$5,000.

We create four data sets:

1. sales of owner-occupied properties,
2. sales of renter-occupied properties,
3. rents of renter-occupied properties, and
4. matched property-level price-rent ratios for renter-occupied properties.

Among these data sets, (1) is a sample of properties distinct from (2), (3), and (4). The properties in (2) and (3) can overlap but do not match exactly, as there are many properties for which we observe a rental listing but do not observe a sale transaction. There are also a few properties sold as investment properties for which we do not directly observe any information on rental income. These are included in (2) but not in (3). Any property that we ever observe being listed for rent or for sale as an investment property is excluded from (1). Data set (1) is a 25 percent random sample of properties in our full data set. Data sets (2), (3), and (4) are 100 percent samples.

The matched sample—data set (4)—comes from two sources. The first is a match of rental listings to sale transactions. We consider any sale within one year (on either end) of the rental listing to be a match. The second source is listings of investment properties that include the net operating income in the listing. We exclude any properties that have a sale price less than their annual rental income and any for which the sale price is greater than 40 times the annual rental income.

Our method of identifying rentals leads to Type I and Type II errors. Some rental properties may never be listed by the MLS or are not identified in for-sale listings as investment properties, so we may misclassify some rentals as owner-occupied. However, the number of misclassified rentals is likely to be small relative to the large number (more than 3.3 million transactions) of owner-occupied transactions in our sample. Our method of identifying rentals is conservative, so we view the reverse misclassification—the erroneous classification of owner-occupied properties as rentals—as unlikely.

³Cooperative buildings are multifamily buildings in which multiple people (usually the residents) own a stake in the ownership of the entire building. These are distinct from condominium buildings, in which property owners own a single unit.

⁴We flag properties as having depreciated substantially if, for example, they were sold “as is” or were advertised as a “handyman special” or “fixer upper.”

There are, of course, properties that transition between owner-occupied and renter-occupied. For example, after the GFC, large buy-to-rent investors bought previously owner-occupied properties, often out of foreclosure, and transitioned them to the rental market (Mills et al. 2019). While it is possible that these transitions drive some of the dynamics in our price-price ratio, we do not think they drive any of our main results for three reasons. First, the entrance of large buy-to-rent investors did not occur until 2012, after the majority of the boom-bust cycle of the 2000s. Second, smaller investors, who were active over the entire time period we study, have no impact on house prices, unlike large investors (Mills et al. 2019). Third, most of the properties purchased by large buy-to-rent investors were acquired at foreclosure auctions (Mills et al. 2019). We exclude foreclosure sales due to concerns about property depreciation between sales, so these transactions are excluded from the owner-occupied price index.

Summary statistics for our data samples for properties with repeat transactions are in Table 1. Panel A includes information for our full sample of sales of both owner-occupied properties and renter-occupied properties as well as our full sample of rent transactions. Panel B includes information on our matched sample. Table A.2 in the appendix contains summary statistics for the full sample, including properties for which we observe only one transaction, confirming that the characteristics of properties with multiple transactions are not materially different from those that transact only once.

A potential issue with our rent data is that most rental units are not listed by the MLS. According to the 2021 Rental Housing Finance Survey, just under 12 percent of rental properties nationwide are listed using a Realtor and thus likely to be listed by the MLS. Furthermore, rental units that are listed by the MLS are higher quality than the average rental unit in the United States. Rental units listed by the MLS are higher priced, in newer buildings, and larger compared with those in the American Housing Survey (AHS). By comparison, sale transactions listed by the MLS are representative, closely matching statistics for newly occupied owner-occupied units in the AHS.⁵ We do not consider the non-representative nature of MLS rental properties to be a first-order concern since we are mainly interested in rent growth, not rent levels. Rent growth as estimated from MLS data is very similar, both at the core-based statistical area (CBSA) level and nationally, to rent growth for new tenants estimated from the BLS’s CPI Housing Survey. The housing survey is a carefully constructed, representative random sample of renter-occupied housing units in the United States and is the microdata underlying CPI tenant rent (Adams et al. 2022).⁶

⁵This can be seen in Table A.1 in the appendix, in which we compare rents on properties listed in the MLS during the 1999–2019 period with rents for recently rented market-rate housing units in the 1999–2019 waves of the AHS.

⁶Figure A.1 in the appendix provides a comparison of our measure of rent inflation to the new-tenant rent index based on the CPI Housing Survey as well as two other publicly available measures of rent inflation.

Outside of Table A.1, when we refer to rent information from the MLS, we mean annual rents net of property taxes. The MLS data often include information on property taxes in the listing. In addition, CoreLogic has matched the MLS data with data collected from local tax assessors. Whenever possible, we net out the actual dollar amount of property taxes associated with a given property. We interpolate tax information for properties for which we have some information but are missing information for a specific year. After calculating the property tax rate (as a percentage of annual gross rental income), we winsorize that rate across our entire sample at 1 percent on either end and use the average of this winsorized rate within each county and year to impute taxes for properties for which we do not have any property tax information.

We do not net out other expenses from our measure of rent. While market-level estimates of costs such as utilities exist, the MLS data do not include consistent information on the finer details of rent contracts. Our repeat-transaction analysis effectively assumes that any additional costs are a constant fraction of the total rent net of property taxes.

3 DECOMPOSITION

3.A A Repeat-Transaction Index Approach

We use a repeat-transaction approach. The theory behind this method is that it purges any characteristics of the property that do not change between the two transactions, including local amenities that are not observable to the researcher and property characteristics such as square footage, number of bedrooms, etc. For this reason, we remove any transaction pair between which there was a substantial renovation or depreciation as described in Section 2. In addition, we remove any transaction pair that implies a greater than 50 percent annual price change.

Formally, we can write the difference in the price or rent of a property at time t_2 and the sale price or rent of the same property at time t_1 , where $t_2 > t_1$, as the sum of the change in the corresponding index between the two transactions:

$$p_{it} = \alpha_i + \sum_{-\infty}^t \phi_t \Rightarrow p_i(t_2) - p_i(t_1) = \sum_{t_1+1}^{t_2} \phi_t.$$

Relying on this relationship, we run the following regression:

$$p_i(t_2) - p_i(t_1) = \sum_{t_1+1}^{t_2} \phi_t + \sum_{t_1+1}^{t_2} \phi_t^{\text{own}} \times \mathbf{I}(p_o) + \sum_{t_1+1}^{t_2} \phi_t^{\text{rent}} \times \mathbf{I}(\text{rent}) + \sum_I \text{CBSA}_i + \epsilon_{it}. \quad (3.1)$$

The left-hand-side variable is the log change in the sale price (for owner-occupied and renter-

occupied properties) or rent (for renter-occupied properties) from one transaction to the next. The ϕ_t are year dummy variables, which are equal to one for $t_1 + 1$, t_2 , and every year from $t_1 + 1$ through t_2 . In addition, we include CBSA by type fixed effects, where the type is an indicator for whether the observation is the transaction price of an owner-occupied property, the sale price of a renter-occupied property, or the rent of a renter-occupied property. The omitted property category on the right-hand side is the sale price of renter-occupied housing; the values of ϕ_t^{own} and ϕ_t^{rent} are relative to the sale price of renter-occupied housing.

The property-year residual is ϵ_{it} . Following Case and Shiller (1987), we weight our regressions to account for heteroskedasticity in the error term due to differences in the time between transactions (transactions further apart should have a larger error term). Following in that tradition, we run our regression in three stages. In the second stage, we run a regression of the squared residuals from the first regression on a constant term and the time between sales. The slope coefficient is a measure of how much the variance of the residuals increases with time between sales. Our final stage regression is a repeat of the first stage, with each observation divided by the square root of the fitted value from the second stage. We do not weight the rental transactions because, in practice, these weights matter little for repeat rent indices (Adams et al. 2022).

3.B Relative Price Growth

The top-left panel of Figure 2 displays a plot of our estimates of ϕ_t^{rent} and ϕ_t^{own} using our full sample.

The top-right panel depicts the price and rent indices implied by our coefficient estimates. The ϕ_t^{own} captures the change in the price index of owner-occupied properties relative to the change in the price index of renter-occupied properties. The changes in the price index of renter-occupied properties are captured by the year fixed effects. Therefore, the repeat-sale index in year t is the cumulative sum of the ϕ_t^{own} and year fixed effects up until year t . The rent index is calculated similarly.

While our repeat-transactions approach controls for individual property characteristics, there still may be systemic differences between the properties for which we observe rents and those for which we observe a transaction price. We therefore also create a price-rent ratio using our sample of properties for which we observe both a transaction price and a rent within the span of one year. We create a weighted repeat-transaction price-rent ratio index as described above by using repeat observations of property-level price-rent ratios. The result is depicted in Figure 2c. It shows a pattern very similar to the result using our full sample, implying that there is little reason to be concerned over the selection of the various samples.

4 IMPLICATIONS

We use a simple model to illustrate what we can infer from our empirical results. We focus on five potential drivers of house price and rent movements: interest rate shocks, preference shocks, house price expectation shocks, supply shocks, and what we refer to as “reallocation shocks” (changes in the economy that lead to a redistribution of housing across the population). Each of these five shocks is capable of increasing house prices, but their implications for p_r , p_o , and $rent$ differ, allowing us to assess whether each shock is a plausible explanation for a given price movement in the data.

Our model, described in detail in the appendix, is populated by infinitely lived households and investors. Our model is in the spirit of Poterba (1984) but builds on Foote et al. (2021). Households decide whether to rent or own and, conditional on that decision, how much housing to consume. Investors decide whether to buy and manage property or to earn a wage in the labor market. Households are heterogeneous in their relative preference for owning versus renting, and investors are heterogeneous in their preferences for managing property versus working in the labor market. Households and investors face potentially binding constraints on housing consumption and investment, respectively. We consider cases where investors can convert housing units from owner-occupied to renter-occupied and vice versa without cost and where conversion is impossible. We abstract from taxes and maintenance costs and think of $rent$ as net operating income (that is, net of any costs of managing the property).

Two equations are key to understanding our model. The first is the user cost equation:

$$user_o = p_o(r - \dot{p}_o/p_o), \tag{4.1}$$

where p_o is the price of owner-occupied housing, and r is the interest rate. The o subscript indicates that the property is owner-occupied, and a similar equation obtains for renter-occupied properties with r subscripts.

The user cost transforms the price of a long-lived asset into a flow cost and is often referred to as the rental rate. In deciding whether to buy real estate and how much to buy, investors compare $rent$ and $user_r$, and households compare $rent$ and $user_o$. Our general formulation allows for the possibility that investors must expend labor effort to manage property. In that case, we can think of the investor as receiving a wage equal to $rent - user_r$ for managing property, and the first-order condition for investors is:

$$rent - user_r \equiv rent - p_r(r - \dot{p}_r/p_r) = \frac{MU_l^i}{MU_c^i}, \tag{4.2}$$

where MU_l and MU_c are the marginal utilities of leisure and consumption, respectively.

In this general formulation, Equation (4.2) implies that the price-rent ratio depends on

investor marginal utilities. However, most housing research (for example, Poterba (1984)) focuses on a special case where investors treat housing as a financial asset. In this case, $MU_i^i = 0$, $rent - user_r = 0$, and the price-rent ratio simplifies to the standard Gordon growth formula:

$$\frac{p_r}{rent} = \frac{1}{r - \dot{p}_r/p_r}. \quad (4.3)$$

The second important equation is the household first-order condition. This depends on whether the household is a renter or owner and is given by:

$$rent = \frac{MU_h^r}{MU_c^r} \text{ or } user_o = \frac{MU_h^o}{MU_c^o}, \quad (4.4)$$

where MU_h is the marginal utilities of housing.

If we assume log utility and focus on the balanced growth path of the economy, we can combine Equations (4.3) and (4.4) to generate fundamentals-based versions of the user cost formula (Equation (4.1)):

$$user_r = p_r(\rho - n), \quad (4.5)$$

and the price-rent formula (Equation (4.3)):

$$\frac{p_r}{rent} = \frac{1}{\rho - n}, \quad (4.6)$$

where ρ is the agent subjective discount rate, and $n = \dot{M}U_h/MU_h$ is the growth rate of the marginal utility of housing. Equations (4.5) and (4.6) follow from the fact that along the balanced growth path:

$$r = \rho - \dot{M}U_c/MU_c = \rho + \dot{c}/c = \rho + \dot{y}/y, \quad (4.7)$$

where the second equality follows from the equilibrium condition that $\dot{c}/c = \dot{y}/y$ and

$$\dot{p}_r/p_r = \dot{rent}/rent = \dot{M}U_h/MU_h - \dot{M}U_c/MU_c = n + \dot{y}/y. \quad (4.8)$$

The second equality follows from the first-order condition for renters and the third from the definition of n , log utility, and the equilibrium condition. Note that although MU_c growth drives both r and \dot{p}_r/p_r , it cancels out in Equations (4.5) and (4.6), which implies that expected productivity growth (\dot{y}/y) has no effect on house prices. While this result is specific to log utility, in a more general setting, we cannot even sign the effect of productivity growth on prices: If the elasticity of intertemporal substitution is greater than one, higher productivity growth could lead to lower prices. Intuitively, the net effect of productivity growth on house prices reflects a horse race between the good news of higher future rents

and the bad news of higher future interest rates.

4.A Interest Rate Shocks

Ceteris paribus, Equation (4.3) implies that interest rates should be negatively correlated with the price-rent ratio. However, the true relationship between interest rates and the price-rent ratio is more nuanced precisely because it is impossible to hold all else equal. Specifically, the relationship between the price-rent ratio and interest rates cannot be understood without accounting for expectations about house price growth.

There are two main drivers of the correlation between expectations about house price growth and interest rates. The first is theoretical. As discussed above, productivity growth affects both price expectations and interest rates: A positive productivity shock leads to both higher interest rates and higher price growth. Which effect dominates depends on preference parameters and is theoretically ambiguous.

The second reason is that monetary policy responds to price growth expectations, including those for house prices. When markets are “frothy”—when participants expect large price growth—policymakers raise interest rates, and when markets are depressed, policymakers cut rates. This pattern was particularly relevant after the financial crisis, when the Federal Reserve lowered rates in response to the fallout from the fall in house prices.

The data confirm the ambiguity of the relationship between the price-rent ratio and interest rates, especially the relationship between market frothiness and monetary policy. Figure 3a is a plot of the annual average price-rent ratio against average seven-year swap rates. Holding \dot{p}_r/p_r constant, we would expect to see the points cluster along a downward-sloping line, but since we cannot control for \dot{p}_r/p_r , the data appear to show no coherent pattern.

Careful examination shows that there are patterns in Figure 3a. From 2000 to 2003, lower mortgage rates do appear to be associated with higher price-rent ratios. However, from 2003 to 2018, the pattern reverses, with higher rates now associated with higher price-rent ratios. This coincides with policymakers raising rates during the 2004–2007 period as the boom gathered steam, before subsequently cutting them in 2008, when the market turned and expectations became highly pessimistic. All this reversed again in 2013 as the housing market started to heal and policymakers started raising rates again, although it is not until 2018 that we see the logical low-rate/high price-rent ratio pattern.⁷

⁷We have focused on the seven-year swap rate, but Figure 3a looks similar whether we use the 30-year fixed-rate mortgage rate or the 10-year Treasury rate. Justiniano et al. (2017) propose that the relevant driver of prices is the spread between rates on privately securitized mortgages and Treasury bonds. In unreported regressions, we find that while the spread has the right sign—a lower spread leads to a higher price-rent ratio—the economic effect is extremely small.

4.B Preference Shocks

Suppose that households have Cobb-Douglas preferences and something happens that increases the budget share allocated to housing. In our framework with a constant housing supply, demand for housing goes up, and to restore equilibrium, prices and rents need to rise. However, while this preference shock may dramatically affect prices, it has no effect on the price-rent ratio. In other words, evidence of a preference shock is a price increase driven by increases in rents without an increase in the price-rent ratio.

Per Figure 1b, there is no evidence of a preference shock in the 2000s boom. We estimate that from 2000 to 2006, real rents rose by less than 2 percent as compared to a 37 percent increase in real p_o and a greater than 50 percent increase in real p_r . Then real rents fell by about 6 percent from 2006 to 2011, as compared to a 41 percent decline in real p_o .

After the GFC, rent growth was much stronger. From 2011 to 2020, real rents grew about 12 percent, or 1.3 percent a year, accounting for nearly half the 28 percent increase in real p_o . One potential explanation for this increased growth was that technological advances changed preferences for housing by, for example, making work from home possible in ways it had never been before, a process that likely sped up during the COVID-19 pandemic.

A preference shock is a logical and plausible explanation for the 2020s housing boom. Figure 1b shows that of the 15 percentage point growth in house prices in 2021, about 10 percentage points, or two-thirds, came from nominal rent growth, and fewer than 4 percentage points, or about one-quarter, came from price-rent growth. While our framework does not rule out other explanations, a preference shock is consistent with the results from other research, such as that of Mondragon and Wieland (2022), who find that work-from-home policies had a substantial impact on house price growth.

4.C House Price Expectation Shocks

House price expectations play an important role in many theories of house price booms (Chodorow-Reich et al. 2021). The idea is that households and investors buy houses not just for the consumption flow or rental income, but also in anticipation of future capital gains. In the context of Equation (4.1), increased house price expectations reduce the user cost of housing and increase demand.

According to Equation (4.8), both n and productivity growth can increase expected price growth, but as illustrated by Equation (4.1), if we assume log utility, productivity growth cancels out of the price-rent equation. Researchers have posited increases in population growth and changes in preferences as sources for changes in n . Foote et al. (2021) point out that, holding the housing stock fixed, increased population growth implies that the housing stock *per capita* shrinks faster over time, which in turn implies higher marginal

utility growth. Kaplan et al. (2020) propose that price expectations are driven by variation in the probability of a switch to preferences with a large budget share for housing in the future.

Overall, the theory implies that an increase in the price-rent ratio is a necessary condition for an expectations-driven increase in house prices. This provides insight into the house price booms in our sample. Figure 1b shows that in 2005 and 2021, house prices rose by more than 10 percent. While in both cases concerns were raised in real time about the role of exuberant expectations, our theory says that increased expectations are a plausible explanation only for the 2005 boom. The reason is that in 2021, almost all growth came from increases in real rents and inflation and very little from increases in the price-rent ratio.

If we follow Poterba (1984) and assume that investors view housing as a purely financial asset, then Equation (4.3) implies something much stronger: All variation in the price-rent ratio results from changes in r and \dot{p}_r/p_r . Thus, given that interest rates were rising in 2005, we can conclude that the *only* possible explanation for the boom in 2005 is increased expectations.

An alternative approach is to look directly at house price expectations. Figure 3b features data from the Federal Reserve Bank of New York’s Survey of Consumer Expectations (SCE), which shows that three-year price expectations did not increase much during the 2021 boom, consistent with our interpretation of the price-rent data. Unfortunately, the SCE was started only in 2014, so we cannot use it to also validate our claim that expectations *did* drive the boom of the 2000s.

4.D Supply Shocks

An increase in supply increases equilibrium quantities of housing and lowers the marginal utility of housing, thus reducing *rent*. If we assume that investors view housing as a purely financial asset, then Equation (4.3) holds, and supply shocks should not affect the price-rent ratio.

Figure 3c illustrates the potential role for housing supply in the 21st century. The figure shows a negative correlation between real rent growth and the ratio of households to housing units. Prior to the GFC, the number of housing units grew faster than the number of households, applying downward pressure on the growth of marginal utility and rents. In the post-GFC period, the pattern reversed. Researchers (for example, see Glaeser and Gyourko (2018)) have argued that the pace of home construction in the wake of the GFC was unusually slow, implying that the quantity of housing per capita fell, raising the marginal utility of housing. In contrast to the relationship between rents and per capita housing, there appears to be little relationship between price-rent growth and growth in the size of households per housing unit. This is consistent with Equation (4.3), which implies that the current marginal

utility of housing should have no effect on the price-rent ratio.

4.E Reallocation Shocks

Up to this point, we have considered shocks that affect all participants in the economy equally. However, in recent years, researchers have focused much attention on what we call *reallocation shocks*. Reallocation shocks redistribute housing across households and/or investors, thereby changing the identity or the holdings of the marginal investor and/or the marginal homeowner.

Reallocation shocks emerged as a focus of research during and after the housing boom of the 2000s. Concurrent increases in house prices and mortgage debt led many to argue that relaxed lending standards were driving the boom. The hypothesis was that since households with good credit did not face binding constraints in credit markets, the relaxation in lending standards gave new buying power to marginal households and therefore led to a reallocation of housing. This argument is often attributed to an influential paper by Mian and Sufi (2009).

Whether or how reallocation shocks affect house prices and rents depends on what one assumes about the economy. We focus here on three key modeling choices: heterogeneity in beliefs, substitutability of owner- and renter-occupied property, and whether investors treat housing as a financial asset.

In what we will call *standard models*, researchers assume homogeneous beliefs and that investors treat housing as a financial asset. Examples of reallocation shocks in standard models include Greenwald (2018), Kaplan et al. (2016), and Favilukis et al. (2017). Equation (4.3) shows that in standard models, the price-rent ratio depends on only r and \dot{p}_r/p_r . Since both are independent of the distribution of housing along the balanced growth path, reallocation cannot affect the price-rent ratio in a standard model.

Whether the reallocation shock affects rents or the price-price ratio depends on whether renter- and owner-occupied properties are substitutes. If they are perfect substitutes, then $p_o = p_r$, and *rent* and prices grow proportionally. This was clearly not the case in the 2000s boom, as real rent growth was exceptionally slow from 2000 to 2006, as discussed in Section 4.B. If we assume no substitutability, then increases in p_o reflect some combination of increases in *rent* and the price-price ratio. This was also not the case in the 2000s boom as the price-price ratio *fell* by almost 20 percent.

Thus, in standard models at least, we can resolve the conceptual issue presented at the 2005 FOMC meeting referenced in the introduction. Given this fall in the price-price ratio, it is unlikely that increased demand for owner- versus renter-occupied property explains the increase in the price of owner-occupied real estate. Notably, reallocation appears to be a more plausible explanation for the 2021 boom, when *rent* and the price-price ratio rose (and

the price-rent ratio did not), accounting for most of the growth.

Equation (4.2) illustrates how reallocation shocks can affect $p_r/rent$ if we relax the financial asset assumption. For example, suppose a reallocation shock results in increased homeownership and reduces the share of property allocated to landlords. Landlords have more leisure, and MU_l^i goes down. To restore equilibrium, rent must fall relative to price, raising $p_r/rent$. Greenwald and Guren (2021) and Sommer et al. (2013) show how reallocation can increase the price-rent ratio in models along similar lines.

But even when we relax the financial asset assumption, explaining the boom without perturbing beliefs is challenging. The issue is the enormous magnitudes of house price changes. Figure 3d shows a counterfactual exercise in which we track the $user_r$, $user_o$, and $rent$ assuming constant house price beliefs and interest rates. Figure 3d implies a 34 percent increase in $user_o$ and a 54 percent increase in $user_r$ relative to $rent$. Confronted by cost increases of this magnitude, why didn't homeowners and, especially, investors exit the market in droves?

Belief shocks and reallocation shocks are not mutually exclusive. Geanakoplos (2010) studies reallocation shocks in a model with heterogeneous beliefs. In his setup, investors are all constrained but differ in their level of optimism about house prices. Relaxing the constraint reallocates housing demand to more optimistic investors. With a more optimistic marginal investor, the $user_r$ must rise relative to rents to clear the market, and the price-rent ratio must increase. As far as the price-rent ratio, the price-price ratio, and rents go, the implications of Geanakoplos-style reallocation shocks and the expectation shocks discussed in Section 4.C above are the same.

5 CONCLUSION

The results in this paper highlight the need to incorporate data on prices of both owner- and renter-occupied properties and rents in any analysis of potential overvaluation in the housing market. As we have shown, without data on the price of renter-occupied properties, one cannot construct a meaningful price-rent ratio, which may result in mistaking an increase in the price-price ratio for an increase in the price-rent ratio or vice versa.

Since Case and Shiller (1987), consistent, robust house price indices have become widely available and widely used. However, market-based rent indices have emerged only more recently with the Zillow Observed Rent Index (ZORI) and the CoreLogic Single-Family Rent Index (SFRI). To date, though, the only price indices available for renter-occupied property are for multifamily properties, which are not directly comparable to the owner-occupied housing stock.

References

- Adams, Brian, Lara Loewenstein, Hugh Montag, and Randal J. Verbrugge (2022). “Disentangling rent index differences: data, methods, and scope.” Working Paper 22-38, Federal Reserve Bank of Cleveland. doi:10.26509/frbc-wp-202238.
- Bracke, Philippe (2015). “House prices and rents: Microevidence from a matched data set in central London.” *Real Estate Economics*, 43(2), pp. 403–431. doi:10.1111/1540-6229.12062. Publisher: Wiley Online Library.
- Campbell, Sean D, Morris A Davis, Joshua Gallin, and Robert F Martin (2009). “What moves housing markets: A variance decomposition of the rent–price ratio.” *Journal of Urban Economics*, 66(2), pp. 90–102. doi:10.1016/j.jue.2009.06.002. Publisher: Elsevier.
- Carson, Joseph G, David S Johnson, and Charles Steindel (2006). “Housing costs in the CPI: What are we measuring?” *Business Economics*, 41(1), pp. 59–68. doi:10.2145/20060108.
- Case, Karl E and Robert J Shiller (1987). “Prices of single family homes since 1970: New indexes for four cities.” Working paper 2393, National Bureau of Economic Research. doi:10.3386/w2393.
- Case, Karl E and Robert J Shiller (1990). “Forecasting prices and excess returns in the housing market.” *Real Estate Economics*, 18(3), pp. 253–273. doi:10.1111/1540-6229.00521. Publisher: Wiley Online Library.
- Chodorow-Reich, Gabriel, Adam M Guren, and Timothy J. McQuade (2021). “The 2000s housing cycle with 2020 hindsight: A neo-Kindlebergerian view.” Working paper 29140, National Bureau of Economic Research. doi:10.3386/w29140. URL <http://www.nber.org/papers/w29140>.
- Davis, Morris A., Andreas Lehnert, and Robert F. Martin (2008). “The rent-price ratio for the aggregate stock of owner-occupied housing.” *Review of Income and Wealth*, 54(2), pp. 279–284. doi:10.1111/j.1475-4991.2008.00274.x.
- Favilukis, Jack, Sydney C. Ludvigson, and Stijn Van Nieuwerburgh (2017). “The macroeconomic effects of housing wealth, housing finance, and limited risk sharing in general equilibrium.” *Journal of Political Economy*, 125(1), pp. 140–223. doi:10.1086/689606.
- Foote, Christopher L, Lara Loewenstein, and Paul S Willen (2021). “Cross-sectional patterns of mortgage debt during the housing boom: evidence and implications.” *The Review of Economic Studies*, 88(1), pp. 229–259. doi:10.1093/restud/rdaa034.

- Geanakoplos, John (2010). “The leverage cycle.” *NBER Macroeconomics Annual*, 24(1), pp. 1–66. doi:10.1086/648285.
- Glaeser, Edward L. and Joseph Gyourko (2008). “Arbitrage in housing markets.” Discussion paper 2151, Harvard Institute of Economic Research,.
- Glaeser, Edward L. and Joseph Gyourko (2018). “The economic implications of housing supply.” *Journal of Economic Perspectives*, 32(1), pp. 3–30. doi:10.1257/jep.32.1.3.
- Greenwald, Daniel L. (2018). “The mortgage credit channel of macroeconomic transmission.” URL <https://gcfp.mit.edu/wp-content/uploads/2016/09/Greenwald-Slides.pdf>.
- Greenwald, Daniel L. and Adam Guren (2021). “Do credit conditions move house prices?” Working paper 29391, National Bureau of Economic Research. doi:10.3386/w29391.
- Hornstein, Andreas (2009). “Problems for a fundamental theory of house prices.” *FRB Richmond Economic Quarterly*, 95(1), pp. 1–24. URL <https://www.fedinprint.org/item/fedreq/26298>.
- Justiniano, Alejandro, Giorgio E Primiceri, and Andrea Tambalotti (2017). “The mortgage rate conundrum.” Working paper 23874, National Bureau of Economic Research. doi:10.3386/w23874.
- Kaplan, Greg, Kurt Mitman, and Giovanni L. Violante (2016). “Consumption and house prices in the Great Recession.”
- Kaplan, Greg, Kurt Mitman, and Giovanni L Violante (2020). “The housing boom and bust: Model meets evidence.” *Journal of Political Economy*, 128(9), pp. 3285–3345. doi:10.1086/708816.
- Mankiw, N Gregory and David N Weil (1989). “The baby boom, the baby bust, and the housing market.” *Regional Science and Urban Economics*, 19(2), pp. 235–258. doi:10.1016/0166-0462(89)90005-7.
- Mian, Atif and Amir Sufi (2009). “The consequences of mortgage credit expansion: Evidence from the US mortgage default crisis.” *The Quarterly Journal of Economics*, 124(4), pp. 1449–1496. doi:10.1162/qjec.2009.124.4.1449.
- Mills, James, Raven Molloy, and Rebecca Zarutskie (2019). “Large-scale buy-to-rent investors in the single-family housing market: The emergence of a new asset class.” *Real Estate Economics*, 47(2), pp. 399–430. doi:10.1111/1540-6229.12189.

- Mondragon, John A and Johannes Wieland (2022). “Housing demand and remote work.” Technical report, National Bureau of Economic Research.
- Pancak, Katherine A (2017). “Variation in local house Price–Rent ratios.” *Journal of Housing Research*, 26(1), pp. 79–94. doi:10.1080/10835547.2017.12092129.
- Poterba, James M (1984). “Tax subsidies to owner-occupied housing: an asset-market approach.” *The Quarterly Journal of Economics*, 99(4), pp. 729–752. doi:10.2307/1883123.
- Smith, Margaret Hwang and Gary Smith (2006). “Bubble, bubble, where’s the housing bubble?” *Brookings Papers on Economic Activity*, 2006(1), pp. 1–67. doi:10.1353/eca.2006.0019.
- Sommer, Kamila, Paul Sullivan, and Randal Verbrugge (2013). “The equilibrium effect of fundamentals on house prices and rents.” *Journal of Monetary Economics*, 60(7), pp. 854–870. doi:10.1016/j.jmoneco.2013.04.017.
- Verbrugge, Randal and Robert Poole (2010). “Explaining the rent–OER Inflation divergence, 1999–2007.” *Real Estate Economics*, 38(4), pp. 633–657. doi:10.1111/j.1540-6229.2010.00278.x.

A. Full Sample


	Transactions	Properties	Price or Rent (\$)	Price or Rent/ Sq. Ft.	Sq. Ft.	Year Built
Sales of Owner-Occupied	3,301,174	1,433,406	309,829	164	1,903	1980
Sales of Renter-Occupied	163,234	75,671	333,834	200	2,508	1961
Rents of Renter-Occupied	3,701,978	1,330,513	18,024	12	1,722	1984

B. Matched Sample

Transactions	Properties	Rent (\$)	Rent/ Sq. Ft.	Price (\$)	Price/ Sq. Ft.	Price/ Rent	Sq. Ft.	Year Built
156,399	75,811	24,491	16	333,503	200	14	2,512	1961

Table 1. SUMMARY STATISTICS. *Note:* These tables exclude properties for which we observe only one relevant transaction. Summary statistics for all transactions are in Table A.2 in the appendix. All dollar values are in 2012 dollars. Not all values are fully populated at the property level. All summary statistics about property characteristics are unweighted averages over transactions (not unique properties). The rental income reported is net of property taxes. The sample of owner-occupied sales is a 25 percent random sample of properties for which we have data. *Source:* Authors' calculations using the CoreLogic Real Estate Database.

(a) Real Prices and Rents



(b) Decomposition of the Price of Owner-Occupied Housing, 2001–2021




Figure 1. THE PRICE-RENT RATIO. *Note:* The top panel shows the Case-Shiller non-seasonally adjusted index and the rent of primary residence series from the consumer price index (CPI). Both are deflated using the CPI excluding food, energy, and shelter. The bottom panel uses our estimated values to decompose the growth of the price of owner-occupied housing into inflation, real rent growth, change in the ratio of the price of owner- versus renter-occupied housing, and growth in the price-rent ratio. *Source:* Authors' calculations using CoreLogic Real Estate Database, CPI components from the BLS, and the CoreLogic Case-Shiller house price index.




Figure 2. RESULTS OF REPEAT-TRANSACTION REGRESSIONS. *Note:* The top two panels depict the results of estimating Equation (3.1). The left panel shows, respectively, the coefficients on owner-occupied and rents of renter-occupied properties, which reflect price growth relative to the prices of renter-occupied properties. The right panel shows the price indices implied by those coefficients. The bottom panel is a repeat-transaction index estimated on repeat observations of property-level price-rent ratios. *Source:* Authors' calculations using CoreLogic Real Estate Database.




Figure 3. IMPLICATIONS OF DECOMPOSING CHANGES IN P_o . *Note:* Figure 3b is based on the NY Fed Survey of Consumer Expectations and the FHFA house price index. Figure 3a is a scatter plot of our annual estimates of $P_r/Rent$ against the average seven-year swap rate. Figure 3c is a graph of the ratio of households to housing units against real rent calculated as our nominal rent index deflated using CPI exclusive of energy, food, and shelter. Figure 3d shows our calculations of the $P_r/Rent$ and $P_o/Rent$ ratios calculated as described in Section 3. *Source:* Authors' calculations using the CoreLogic Real Estate Database, CPS, Census, NY Fed Survey of Consumer Expectations, FHFA, BLS, and Haver Analytics.

A APPENDIX

A.A Supplemental Figures and Tables

This section contains supplemental figures and tables referenced in the text.

	Renter-Occupied			Owner-Occupied	
	AHS	MLS Unweighted	MLS Weighted	AHS	MLS Unweighted
Characteristics					
Rent or Price (2010 \$)	919	2,882	3,855	225,662	272,820
Year Built	1967	1977	1971	1976	1977
Bedrooms (#)	2	3	2	3	3
Bathrooms (#)	2	2	2	3	2
Size					
Share < 500 Sq. Ft.	7	1	7	1	0
Share 500–750 Sq. Ft.	19	4	19	2	2
Share 750–1,000 Sq. Ft.	27	10	27	7	7
Share 1,000–1,500 Sq. Ft.	29	32	30	23	31
Share 1500+ Sq. Ft.	17	53	17	61	60

Table A.1. COMPARISON OF AHS AND MLS. *Note:* Values for the AHS are weighted averages from the 1999–2019 surveys and are limited to households that had moved since the preceding survey. Rental units from the AHS exclude all rent-controlled and subsidized housing units. Values from MLS are from listings closed in 1999–2019. *Source:* AHS and CoreLogic MLS.

A. Full Sample

	Transactions	Properties	Price or Rent (\$)	Price or Rent/ Sq. Ft.	Sq. Ft.	Year Built
Sales of Owner-Occupied	6,943,246	5,076,965	312,827	165	1,911	1979
Sales of Renter-Occupied	1,093,637	1,006,074	321,622	215	2,117	1973
Rents of Renter-Occupied	5,416,945	3,175,621	22,449	16	1,718	1982

B. Matched Sample


Transactions	Properties	Rent (\$)	Rent/ Sq. Ft.	Price (\$)	Price/ Sq. Ft.	Price/ Rent	Sq. Ft.	Year Built
1,030,623	1,006,522	22,036	16	312,649	215	14	2,118	1973

Table A.2. SUMMARY STATISTICS FOR SAMPLE INCLUDING PROPERTIES THAT TRANSACT ONLY ONCE. *Note:* Unlike Table 1, these tables *include* properties for which we observe only one relevant transaction. All dollar values are in 2012 dollars. Not all values are fully populated at the property level. All summary statistics about property characteristics are unweighted averages over transactions (not unique properties). The rental income reported is net of property taxes. The sample of owner-occupied sales is a 25 percent random sample of properties for which we have data. *Source:* Authors' calculations using the CoreLogic Real Estate Database.

Year	P_r	P_o	Rent	$\frac{P_r}{Rent}$	$\frac{P_o}{P_i}$	$\frac{P_o}{Rent}$
2000	0.00	0.00	0.00	0.00	0.00	0.00
2001	0.10	0.07	0.04	0.07	-0.03	0.03
2002	0.23	0.13	0.05	0.18	-0.10	0.08
2003	0.36	0.20	0.06	0.30	-0.15	0.15
2004	0.49	0.30	0.07	0.42	-0.18	0.23
2005	0.63	0.42	0.08	0.55	-0.21	0.34
2006	0.66	0.46	0.11	0.55	-0.20	0.35
2007	0.59	0.41	0.14	0.45	-0.18	0.27
2008	0.43	0.30	0.11	0.32	-0.13	0.18
2009	0.26	0.22	0.08	0.18	-0.04	0.14
2010	0.17	0.18	0.11	0.06	0.02	0.08
2011	0.14	0.13	0.14	-0.00	-0.01	-0.01
2012	0.20	0.14	0.18	0.02	-0.06	-0.04
2013	0.31	0.21	0.21	0.09	-0.10	-0.00
2014	0.36	0.25	0.23	0.13	-0.11	0.02
2015	0.42	0.29	0.26	0.15	-0.13	0.03
2016	0.47	0.34	0.29	0.18	-0.13	0.05
2017	0.54	0.39	0.30	0.24	-0.15	0.09
2018	0.60	0.44	0.33	0.28	-0.17	0.11
2019	0.65	0.47	0.35	0.30	-0.18	0.12
2020	0.72	0.53	0.38	0.34	-0.19	0.15
2021	0.85	0.68	0.48	0.37	-0.17	0.20

Table A.3. INDICES FROM FIGURE 2B. *Note:* All indices are log points relative to 2000. *Source:* Authors' calculations using the CoreLogic Real Estate Database.

Annual Rent Inflation



Rent Inflation Measures over Subperiods




Figure A.1. COMPARING RENT INFLATION MEASURES. *Note:* These graphs compare annual new-tenant rent inflation from a variety of sources. All of these sources use a repeat-transaction methodology. The MLS Rent Index is our index. NTRR is from Adams et al. (2022) and is created using the same microdata that underlie CPI tenant rent, which is a geographically representative random sample of renter-occupied housing units in the United States. SFRI is the CoreLogic Single-Family Rent Index, which is also based on the MLS but is restricted to single-family homes and is created using a weighted average of city-level indices. ZORI is the Zillow Observed Rent Index, which is based on proprietary data. The bottom panel plots annual average inflation for the specified groups of years.

B MODEL

Our baseline model is a continuous-time, nonstochastic endowment economy with a fixed stock of housing based on the model in Foote et al. (2021). Housing should be thought of as land, which does not depreciate. We assume the existence of purchase and rental markets, and our setup is flexible enough to handle the case where properties are perfectly substitutable across tenure types and where markets are perfectly segmented. We focus on the user cost in the spirit of Poterba (1984), but our treatment more closely follows Hornstein (2009). Our economy is composed of infinitely lived households and investors and a finite supply of property.

B.A Households

Households receive labor income, can borrow and save, can allocate their expenditure across a consumption good and housing, and choose whether to rent or own housing. Households maximize their lifetime utility:

$$\max_{\{c_t, h_t, I_t\}_{t=0}^{\infty}} \int_{t=0}^{\infty} e^{-\rho t} N_t u(c_t, h_t/N_t, I_t) dt,$$

where ρ is the subjective discount rate; N is household size, which grows at rate $\dot{N}/N = n$; h is real units of housing; c is consumption; and I is an indicator that equals one if the household owns its home and zero if it rents.

If a household opts to own, it allocates its flow income to saving, consumption, and net additions $x_{o,t}$ to their stock of housing at unit price p_t^o , the price of owner-occupied housing. There is a single financial asset that pays an endogenous interest rate r_t , meaning that homeowner financial wealth a_t grows at:

$$\dot{a}_t = N_t y_t + r_t a_t - N_t c_t - p_t^o x_t^o, \tag{B.1}$$

and holdings of housing grow at $\dot{h}_t = x_t^o$. The dynamic budget constraint (Equation (B.1)) and a no-ponzi-game condition yield the lifetime budget constraint:

$$\int_{s=t}^{\infty} e^{-R(t,s)} (N_s c_s + p_s^o (r_s - \dot{p}_s^o/p_s^o) h_s) ds = a_t + p_t^o h_t + \int_{s=t}^{\infty} e^{-R(t,s)} N_s y_s ds. \tag{B.2}$$

Equation (4.1) follows from the fact that the relevant price of housing in Equation (B.2) is not p_t but rather the user cost $user_t^o = p_t^o (r_t - \dot{p}_t^o/p_t^o)$.

If a household rents, it allocates its flow income to consumption, rent, and savings, so its

wealth grows at:

$$\dot{a}_t = N_t y_t + r_t a_t - N_t c_t - rent_t h_t,$$

which yields a lifetime budget constraint identical to Equation (B.2) except that rent replaces the user cost.

Maximizing household utility subject to Equation (B.2) yields Equation (4.4) for owners and, with rent substituted for user cost, for renters.

We capture the idea of credit constraints by introducing limits on how much of their budget households and investors can spend on housing. For renters, this constraint takes the form $rent \cdot h \leq \theta^r y$, and for owners, $user^o \cdot h \leq \theta^o y$.

B.B Investors

Investors maximize lifetime utility

$$\max_{\{c_t, l_t, I_t\}_{t=0}^{\infty}} \int_{t=0}^{\infty} e^{-\rho t} N_t u(c_t, l_t, I_t) dt,$$

where l_t is an investor's leisure time, and I is an indicator equal to one if the investor chooses to invest in real estate and zero otherwise. We assume that an investor can manage λ units of housing per unit of labor, so $h = (L - l)/\lambda$ is the rental housing supplied by a given investor. Investor flow utility is

$$u(c_t, l_t, I_t) = \log(c_t^\eta l_t^{1-\eta} \omega^I) dt.$$

If the investor chooses to be an investor, wealth accumulation equals:

$$\dot{a}_t = rent_t h_t + r_t a_t - N_t c_t - p_{r,t} x_{r,t},$$

and the investor sets $\dot{h}_t = x_{r,t}$ using the identity that $h = (L - l)/(\lambda N)$.

$$\int_{s=t}^{\infty} e^{-R(t,s)} N_s \left(c_s + \frac{rent_s - uc_s}{\lambda N_s} l_s \right) ds = a_t + p_t h_t + \int_{s=t}^{\infty} e^{-R(t,s)} N_s \frac{rent_s - uc_s}{\lambda N_s} L_s ds. \quad (\text{B.3})$$

Maximization of investor utility subject to Equation (B.3) yields Equation (4.2).

If the investor chooses to be a worker, they earn $wage/N$ per unit of labor supplied:

$$\dot{a}_t = wage_t (L_t - l_t) + r_t a_t - N_t c_t.$$

So

$$\int_{s=t}^{\infty} e^{-R(t,s)} (N_s c_s + wage_s l_s) ds = a_t + \int_{s=t}^{\infty} e^{-R(t,s)} wage_s L_s ds. \quad (\text{B.4})$$

B.C Balanced Growth

If we assume that households have Cobb-Douglas utility across goods and housing and log utility over time, then we can characterize the balanced growth path. All participants in the market in the model must satisfy the same Euler equation:

$$\dot{u}_c/u_c = \rho - r \Rightarrow r - \rho = \dot{c}/c = \dot{y}/y. \quad (\text{B.5})$$

Price growth must equal rent growth:

$$\dot{p}/p = \dot{rent}/rent = \dot{c}/c + n = \dot{y}/y + n,$$

where the second equality follows from taking derivatives of Equation (4.4).

B.D Effect of Expectations

We define an expectations-driven boom as an increase in beliefs about the growth rate of the marginal utility of housing (n), which in turn implies an increase in \dot{p}_r/p_r . An increase in n to n' implies an increase in the price-rent ratio because p_r increases to $p'_r = p_r(\rho - n)/(\rho - n')$, but the $rent$ is unchanged because, while expectations about the future have changed, the current demand for housing is unchanged, and $rent$ is effectively the spot price of housing.