The Riskiness of Owning Versus Renting Housing*

Lee M. Lockwood† Scott R. Baker‡ Lorenz Kueng§ Pinchuan Ong¶

Preliminary; please don't cite or circulate without permission

Click here for most recent version

September 11, 2019

Abstract

Homeowners and renters have mirror-image exposures to the substantial risks in housing prices. The costs of these exposures depend crucially on their correlations with other important exposures in household portfolios. Using over 70 years of data on local markets in the U.S., we find that rents and home prices are strongly positively correlated with wages at all horizons. As a result, renting insures earnings risk, and—contrary to widely-held views—for many households owning is much riskier than renting. Combined with evidence that many households view owning as being particularly safe, our findings suggest that the efficiency costs of the substantial tax preferences for owner-occupied housing are greater than previously thought.

*We thank Jeff Campbell, Eric Chyn, Kerem Cosar, Leland Farmer, Greg Howard, Matt Notowidigdo, Ed Olsen, Isaac Sorkin, and seminar participants at Northwestern, UVA, IFS (UCL/LSE), NYU, Claremont McKenna, the Chicago Fed, Geneva, Luzern, Munich, Salzburg, William & Mary, and the Empirical Macro Workshop Las Vegas for helpful discussions.

†First author. University of Virginia, Department of Economics, 248 McCormick Road, Charlottesville, VA 22903, and National Bureau of Economic Research. E-mail: leelockwood@virginia.edu.

‡Northwestern University, Department of Finance, 2001 Sheridan Road, Evanston, IL 60208, and National Bureau of Economic Research. E-mail: s-baker@kellogg.northwestern.edu.

§University of Lugano (USI), Department of Economics, Via Buffi 13, Lugano, TI 6900; Northwestern University, Department of Finance; Centre for Economic Policy Research; and National Bureau of Economic Research. E-mail: lorenz.kueng@usi.ch.

¶Northwestern University, Department of Economics. E-mail: PinchuanOng2014@u.northwestern.edu.

1 Introduction

Housing is a major consumption good and an important asset, comprising about one third of both household expenditures and household assets. A household's exposure to the substantial risks in rents and home prices depends crucially on whether it owns or rents its home and for how long. Owners are exposed to risk in the price at which they will sell their home, whereas renters are exposed to risk in the rents they will pay for their housing. Whereas a longer horizon in a home increases a renter's exposure to rent risk by accumulating more draws of risky rent, it eventually decreases an owner's exposure to home price risk by discounting the risky sale price farther into the future. A household with a long enough horizon in a home—on the order of six years or more—therefore faces less risk in its present value of housing costs by owning than by renting (Sinai and Souleles, 2003). As a result, the prevailing view is that for many households owning is safer than renting.[1]

In this paper, we highlight the crucial role of correlations between housing and other items in households' portfolios in determining the riskiness of owning versus renting housing. Although owning tends to lead to less risk in the *housing costs* of households with long enough horizons in a home, renting tends to lead to less risk in their total portfolios—and so in their consumption and welfare. The difference arises because of strong, opposite-signed correlations between the housing exposures of owners versus renters and wages. Whereas the value of an owner's housing exposure is strongly positively correlated with wages, the value of a renter's housing exposure is strongly negatively correlated with wages: Higher wages are strongly associated with higher home prices and rents. Hence, owning exacerbates the substantial wage risk that many households face, whereas renting hedges it. As a result, for many households owning is much riskier than renting, regardless of their horizon in their home.

We begin by using over 70 years of data on local markets in the U.S. (commuting zones) to analyze the evolution of wages, rents, and home prices over life cycle-relevant frequencies. This analysis produces two key findings. First, there is substantial variation both in wages, rents, and home prices over time and across locations, and in the growth of wages, rents, and home prices across different locations during a given period of time. The variation is sufficiently large and persistent to have a major impact on a household's lifetime budget constraint. This suggests that risk in wages, rents, and home prices is a major driver of risk in lifetime consumption. Second, home prices and rents are strongly positively correlated with wages. This is true not only in the cross-section at a point in time and on average over

[1] The view that owning is safer is widespread in the economics literature (e.g., Sinai and Souleles, 2005), in popular financial advice (e.g., Malkiel, 1999), and among U.S. households, even in the wake of the housing bust of the Great Recession (Adelino, Schoar and Severino, 2018).

time but, crucially, in the cross-section of growth rates over time across locations. That is, locations that experience faster growth in home prices and rents during a particular period of time are much more likely to have experienced faster growth in wages as well, over all horizons from 10 to 70 years. This suggests that the substantial wage risk to which many households are exposed would be exacerbated by owning but partially hedged by renting.

Using a variety of approaches, we investigate the implications of these strong empirical regularities for the riskiness of owning versus renting housing. These approaches range from calculations of the net income (non-housing consumption) of hypothetical households who own versus rent in different locations and time periods to calculations of the welfare cost of different housing exposures in a rich life cycle model. All point to the same conclusion: For most households, renting is safer than owning, usually significantly so. For example, we estimate that even for an otherwise-typical household that never moved, the reduction in lifetime consumption risk from switching from always owning to always renting would be worth over 1.5 percent of lifetime consumption, or about $30,000.

The sizable risk benefit of renting over owning for many households reflects two reinforcing factors, both of which are increased by the strong correlations between housing prices and wages. First, renters avoid the substantial sales price risk facing owners. Home prices are highly variable—much more so than rents—and this risk is made more costly by the strong correlation between home prices and wages for two reasons. First, this correlation means that an owner's payoff is negatively related to marginal utility: The payoff from owning tends to be greatest in the high-wage states of the world in which marginal utility is relatively low and smallest in the low-wage states of the world in which marginal utility is relatively high. We estimate that for a typical household, owning amplifies lifetime earnings risk by 30–100% and its welfare cost by even more. Second, this correlation reduces the value of an owner's hedge against risk in future housing costs in different locations. As Sinai and Souleles (2005) emphasize, owning can hedge risk in future housing costs if households tend to move between markets with positively correlated housing prices. The strong correlation between wages and housing prices reduces the value of this hedge by reducing the cost of risk in future housing costs to working households, since higher housing costs are usually compensated by higher wages.

The second, more important factor driving the sizable risk benefit of renting over owning for many households is that renting provides a valuable hedge against the substantial earnings risk they face from risk in market-level wages. We estimate that renting hedges about 30% of the lifetime earnings risk from market-level wage risk faced by a typical household. Many households' exposure to wage risk is sufficiently costly that the benefit of hedging it by renting exceeds the direct cost of being exposed to rent risk, and so exposure to rent risk *reduces* overall risk exposure. That is, for a typical working household, renting is not only safer than owning, it is safer than being fully insured against risk in housing prices. For such households, that owning hedges rent risk is not a benefit of owning but a cost. By owning, a household not only exposes itself to sales price risk, it also forgoes the valuable hedge against earnings risk that renting would provide.

The crucial role of a household's exposure to wage risk highlights systematic heterogeneity in the risk consequences of owning versus renting housing. Single-earner households enjoy a smaller risk benefit from renting than dual-earner households do, and retirees enjoy a benefit that is smaller still; in fact, for many retirees owning would be safer. Hence, the risk-minimizing life cycle profile involves renting during prime working years and owning during retirement. Households in locations with more volatile labor demand and less elastic housing supply enjoy a larger risk benefit from renting, since in such locations wages are both more volatile and more strongly correlated with rents and home prices.

Our main contribution is to a better understanding of the riskiness of owning versus renting housing. Although several papers consider how wage risk might affect housing tenure choices, we go beyond existing work by using long-run historical evidence to quantify the welfare costs of different housing-related exposures in a variety of ways and for different types of households.[2] Our central conclusion that for many households renting is safer than owning contrasts with the prevailing view in the research literature, in financial advice, and among households. Our findings highlight important dimensions of heterogeneity, many of which also differ from prevailing views. Although the riskiness of owning versus renting is just one factor among many in determining the best choice for any given household, our results suggest that it is a quantitatively important factor that is widely misunderstood.

Our findings also have important implications for government policies. That for many households owning is riskier than renting does not, by itself, have immediate policy implications. But together with evidence that most households view owning housing as being extremely safe (Adelino et al., 2018), our results suggest that many households underestimate the relative riskiness of owning. Such misperceptions would magnify the welfare cost of the substantial net subsidies to owner-occupied over rental housing. Likewise, the subsidies would magnify the welfare cost of such misperceptions. Our back-of-the-envelope calculations suggest that misperceptions about housing risks could increase the efficiency cost of

[2]For example, Campbell and Cocco (2001), Cocco (2000), and Yao and Zhang (2001) investigate optimal housing tenure choice in quantitative life cycle models with correlated wages and house prices. Davidoff (2006) investigates how heterogeneity in the correlation between a household's earnings and home prices—due to heterogeneity in the industries in which they work and the MSAs in which they live—relates to heterogeneity in the amount of housing people own. We depart from these papers in our focus on the welfare cost of different housing exposures, which we investigate by analyzing the evolution of wages, rents, and home prices over a sufficiently long time period to capture life cycle-relevant dynamics.

housing subsidies by over $5,000 per household over its lifetime and reduce welfare by even more.

Our paper also contributes to a better understanding of labor earnings risk. We provide new evidence on labor earnings risk over life cycle-relevant frequencies and estimate its welfare cost for different households. This complements and extends the large literature on labor earnings risk, much of which focuses on characterizing individual- or household-level idiosyncratic risk over shorter time horizons (see Karahan, Guvenen, Ozkan and Song (2015) for a recent review). Although idiosyncratic risk drives much of the risk in a household's earnings over relatively short horizons, market-level risk drives much of the risk in a household's earnings over its lifetime. Market-level risk in turn is driven in large part by location-specific risk: In any given time period, some locations experience much faster wage growth than others, and these differences are sufficiently persistent to generate substantial risk in lifetime earnings through this channel alone. Our findings also help clarify the welfare and behavioral effects of wage risk by highlighting the crucial role of a household's housing tenure status in determining its full exposure to wage risk.

2 Theory: Housing Tenure Choice and Net Income Risk

Renting and owning involve mirror-image exposures to risk in housing costs. Renters have a negative exposure to rents during their stay in a home, whereas owners have a positive exposure to expected future rents after their stay through their sale price. Owners effectively prepay future rents, both rents during their stay and, due to the durability of housing, rents after their stay as well. Compared to renting, owning therefore eliminates the negative exposure to rents during the stay but creates a positive exposure to expected future rents at the time of sale.

Because market frictions and other factors limit trade in state-contingent claims, a household's decision about whether to rent or own its home affects its overall risk exposure. The welfare cost of a given exposure depends on its relationship with the rest of the household's portfolio. The most important asset for many households, and one whose risk is difficult to diversify away, is human capital—labor earnings.

The conceptual thought experiment we focus on in this paper is of a household facing exogenous risk in its earnings and housing costs from risk in the wages, rents, and home prices in its labor and housing markets. In reality, households can re-optimize in response to price changes, which makes price risk different from net income risk. Ignoring the effects of re-optimization or, equivalently, assuming perfectly inelastic labor supply and housing demand with respect to changes in these prices, is a convenient approximation that captures the first-order welfare effect of price changes.[3] This experiment focuses the analysis on budget constraints. The key question is: What does the risk in wages and housing prices mean for the risk in net income (income net of housing costs)—and so living standards—of renters and owners?

Consider the implications of wage and housing price risk for risk in the present value of a household's net income during its stay in a home, from period 1 to T. The present value of net income, C, equals the present value of income, Y, less the present value of net housing costs, H: $C = Y - H$. The present value of net housing costs is

$$H = \begin{cases} \sum_{t=1}^{T} \frac{R_t}{(1+r)^{t-1}} & \equiv R \quad \text{if rent,} \\ P_0 - \frac{P_t}{(1+r)^T} & \equiv O \quad \text{if own,} \end{cases}$$

where R_t is the rent in period t, P_t is the home price in period t, and r is the real interest rate.[4]

Net income risk, measured as the variance in its present value, is

$$V(C) = V(Y) + V(H) - 2Cov(Y, H). \tag{1}$$

The riskiness of a given housing exposure depends not only on its riskiness as viewed in isolation, $V(H)$, but also on its covariance with the rest of the household's portfolio, $Cov(Y, H)$. Other things equal, housing costs that covary positively with income are less costly because they reduce the welfare cost of income risk.

Start by considering the special case in which housing costs are uncorrelated with the rest of the portfolio, $Cov(Y, H) = 0$. In this case, the variance of net income is $V(C) = V(Y) + V(H)$, and the riskiness of owning versus renting amounts to a comparison of the variance of the present value of rents during the household's stay to the variance of the

[3] The mechanical effect of a price change on a household's budget constraint—i.e., the implied change in its income net of expenditures holding its behavior fixed—is, by the envelope theorem, a first-order approximation to the full welfare effect of the price change. The full welfare effect includes higher-order terms in the size of the price change, which reflect the gain from re-optimization in response to the change. The quality of the approximation is better the less elastic are labor supply and housing demand with respect to these key prices.

[4] In addition to payments from buying and selling their homes, in reality owners also incur carrying costs such as property taxes and maintenance and repair costs. (Mortgage payments are part of the capital structure decision, not the capital budgeting decision.) Much of the literature on housing tenure choice ignores the risk properties of carrying costs. We exclude them from this analytical section to simplify its exposition, but we allow for them in the quantitative analyses to follow.

present value of the home price at the end of the household's stay:

$$V(C_{rent}) > V(C_{own}) \iff V(R) > V(O) \iff V\left(\sum_{t=1}^{T} \frac{R_t}{(1+r)^{t-1}}\right) > V\left(\frac{P_T}{(1+r)^T}\right).$$

In this case, the relative riskiness of owning versus renting depends crucially on the household's horizon, the length of time it will spend in the home, since lengthening the horizon has opposite-signed effects on the variance of housing costs of owners versus renters. An increase in horizon increases a renter's exposure to rent risk, since it exposes him to rents farther into the future in addition to the rents he was already exposed to. By contrast, under standard assumptions about discounting and home price risk, an increase in horizon eventually decreases an owner's exposure to home price risk since it pushes farther into the future—and so discounts to a greater extent the risk from—the eventual sale of the home. As the horizon tends to infinity, the owner's housing risk approaches zero: $V(O) = V(P_T/(1+r)^T) \to 0$. More generally, for households with sufficiently long horizons, owning is safer than renting: $V(C_{rent}) - V(C_{own}) \to V(R) > 0$.[5] Sinai and Souleles (2003) present evidence that the threshold horizon beyond which a household's present value housing costs are less variable if it owns is not long—on the order of six years. This evidence has played a key role in shaping the view that for many households owning is safer than renting.

If housing costs are correlated with the rest of household's portfolios, however, the relative riskiness of owning versus renting depends not only on the "direct effect"—the variance of housing costs—but also on the "portfolio effect"—the covariance between housing exposures and the rest of household portfolios. The portfolio effect can have a major impact on both the absolute risks of owning and renting housing and on the relative risk of owning versus renting.

Most relevant for housing tenure choice, the portfolio effect systematically affects the relative riskiness of owning versus renting. It is opposite-signed for renters and owners because of their opposite-signed exposures to housing costs. Renters are short housing during their stay; high rent realizations reduce their net income. Owners are long housing after their stay; high expected future rents after their stay increase their sale price and so their net income. The different-signed exposures of renters and owners to housing costs implies different-signed portfolio effects. If housing costs are positively correlated with income, the portfolio effect reduces the riskiness of renting (since rent risk hedges income risk) and increases the riskiness of owning (since home price risk exacerbates income risk).

[5] The role of discounting is crucial, since the durability of housing tends to magnify the effect of rent risk on home price risk. With perfect durability (infinitely-lived homes and land), the increase in current home prices from a shock expected to increase flow rents by ΔR forever is $\Delta R/r$, or about $33\Delta R$ when $r = 3\%$.

The portfolio effect can even change the sign of the contribution of housing risk to net income risk. Without the portfolio effect, exposure to housing risk necessarily increases net income risk. With the portfolio effect, however, exposure to housing risk can potentially decrease net income risk compared to a situation in which housing risk is perfectly insured, i.e., $V(C) < V(Y)$. This occurs if the implicit payoff of the housing position is sufficiently negatively correlated with the rest of the portfolio. For example, exposure to rent risk reduces net income risk if $2Cov(Y, R) > V(R)$ (applying equation (1) to a renter), i.e., if the slope of the regression of income on rents exceeds one half, $\beta_{Y|R} > 1/2$. If rents and income are correlated enough and income risk is large enough relative to rent risk, a household that cannot otherwise hedge its income risk benefits from exposure to rent risk, since rent risk insures income risk. In this case, the fact that owning reduces exposure to rent risk is a cost, not a benefit, of owning.[6]

The portfolio effect also fundamentally alters how changes in horizon affect the relative riskiness of owning versus renting. As discussed above, in the absence of a portfolio effect and under standard assumptions about discounting and home price risk, lengthening the horizon eventually reduces the relative riskiness of owning and makes owning safer. With the portfolio effect, however, lengthening the horizon has an ambiguous effect on the relative riskiness of owning versus renting. If renting provides a valuable hedge against earnings risk, the value of this hedge tends to increase in the horizon. This works against the more widely-appreciated factor that an owner's sales price risk is eventually decreasing in the horizon.

The key point of this analysis is that in addition to the variances of rents and home prices and a household's horizon, a crucial determinant of the relative riskiness of owning versus renting is the relationship between housing costs and the rest of household portfolios. We turn now to estimating the relationships between housing costs and a key determinant of the value of household portfolios: wages.

3 Data: Wages, Rents, and Home Prices

Our main data source is the IPUMS-USA Decennial Census database from 1940 to 2010 (Ruggles, Flood, Goeken, Grover, Meyer, Pacas and Sobek 2018). Its key features for our purposes are its large sample size, detailed geographic and demographic information, and long time coverage. Together, these enable us to construct measures of earnings, rents, and home prices at the local market level over a 70-year time period. The long time coverage in

[6] Owning reduces net income risk relative to the case of no exposure to housing costs if home prices are sufficiently *negatively* correlated with income, $2Cov(Y, P_T/(1+r)^T) < -V(P_T/(1+r)^T) \iff 2Cov(Y, P_T) < -V(P_T)/(1+r)^T$, or $\beta_{Y|P_T/(1+r)^T} < -1/2$.

particular is crucial for learning about the risk facing households over life cycle horizons.[7]

Our aim, following the conceptual thought experiment introduced in Section 2, is to learn about the risk in earnings and housing costs that households face as a result of the risk in the wages, rents, and home prices in their labor and housing markets. Constructing market prices from household-level data involves two main steps: defining markets and measuring the prices in those markets.

Our baseline market definition is a Commuting Zone (CZ). A CZ is a county or set of counties chosen so as to approximate a single labor market with a common workforce. CZs are widely used in economics research as approximations to markets. They also have the advantage of being amenable to long-run analyses, since county boundaries are stable over time, whereas the boundaries of other common definitions of local labor markets (e.g., MSAs) change frequently over time. For households whose county is not uniquely identified, we use IPUMS-supplied allocation factors for each Decennial Census to assign their counties (see Appendix A.1.1). We drop Alaska and Hawaii because their county population data are not available before they obtained statehood in 1959. This leaves us with 721 out of 744 CZs (or 3,060 out of about 3,140 counties) covering 99% of the U.S. residential population.

The second main step is to measure market prices. This involves choosing samples and constructing market prices from individual- and household-level data. Our main sample is households not living in group quarters whose highest earner (i.e., "representative household member") is between 25 and 60 years old. For individual-level variables such as wage income or age, we use the values of the representative household member.[8] To measure wages, we further restrict the sample to households whose representative member worked at least 50 weeks in the past year and at least 35 hours in the previous week or "usual week," depending on the year. These sampling restrictions are aimed at producing cleaner measures of market prices. The age restriction, for example, reduces the effect of schooling and retirement on measured wages.

To construct market prices, we follow the common practice of using hedonic regressions to remove the effects of observable differences in quantities. Separately for each Census year, we regress wage income on fixed effects for household age, and we regress rents and home prices on fixed effects for home age and home size. The CZ median of the estimated residuals are our market price measures.[9] All analyses use household sampling weights. All dollar

[7]To the best of our knowledge, no other source can match the Census's length of coverage of the key variables of interest. Moreover, the quality of the resulting price measures appears to be quite good. For example, in periods of overlap, we find that home prices in the Decennial Census track closely those in home transaction data such as Zillow or FHFA; see Appendix B.1.6.

[8]An alternative would be to use the values of the "head of household," but the definition of the head has changed considerably over this period (e.g., see https://usa.ipums.org/usa-action/variables/RELATE).

[9]Using medians rather than means minimizes the effect of changes in top-coding over time, which could

values are converted to real 2010 dollars using the CPI All Urban series. Full details of the analyses are in Appendix A.

4 Wage, Rent, and Home Price Variation and Correlations

[Table 1 about here]

The first panel of Table 1 shows summary statistics of market wages, rents, and home prices in the 2010 cross section of 721 CZs. Wages are the annual earnings of full-time workers and rents are annual rents. Wages, rents, and home prices exhibit substantial variation across locations. The standard deviations and coefficients of variation are $8,913 and 17.4% for wages, $2,946 and 31.6% for rents, and $129,090 and 54.8% for home prices. Rents and home prices are highly correlated with wages across locations; both correlations are about 0.80. As predicted by spatial equilibrium models, locations with high wages also tend to have high housing costs.

Cross-sectional variation, however, is relatively uninformative about the risk in these key prices. To come closer to a measure of risk, in the remaining five panels of Table 1 we consider the growth rates in these prices over 10-, 40-, and 70-year periods, both overall and relative to the national average. Growth rates are defined as log changes over different horizons k, $\ln x_{c,t+k} - \ln x_{c,t}$, where c is CZ. By focusing on growth rates rather than levels, we isolate the effects of changes over time within locations, setting aside the level differences across locations. This is likely closer to a measure of risk. It follows the common practice of proxying for risk, i.e., the unobservable variation across states of the world, with some subset of the observed cross-sectional and time-series variation.

The results reveal three key patterns. First, variation in the growth rates of these key prices over time and across locations is large and persistent, including across locations during a given period of time. Ten-year log changes of wages and rents each have a standard deviation of about 16 log points and a 90-10 percentile difference of 37 log points. The variation in 10-year changes is not rapidly reversed. Seventy-year log changes of wages and rents have a standard deviation of about 18 log points and a 90-10 percentile difference of about 43 log points. This is suggestive that risk in these key prices might drive significant risk in living standards over the life cycle.

otherwise lead to spurious correlations. The results are similar if we use means instead of medians (once top and bottom codes are made consistent over time) and if we use raw instead of residualized measures. See Appendix B.

Second, home prices are much more variable than rents. Ten-year log changes of home prices have a standard deviation of 27 log points (versus 16 for wages and rents) and a 90-10 difference of 69 log points (versus 37 for wages and rents). This is consistent with the finding of a large literature that home prices are more volatile than their rental cash flows (e.g., Campbell, Davis, Gallin and Martin (2009), Plazzi, Torous and Valkanov (2010)).

Third, rent and home price growth are highly positively correlated with wage growth at all frequencies. For example, the correlation between 40-year log changes in wages and rents is 76% and in wages and home prices is 74%. Over the same horizon, the correlation between *relative* log changes (relative to the national average change) in wages and rents is 69% and in wages and home prices is 48%.

[Figure 1 about here]

To get a better sense of the underlying relationships between these key prices, Figure 1 shows scatter plots of rent growth on wage growth (left panels) and home price growth on wage growth (right panels). Each circle represents a CZ-period and has an area proportional to population. Observations from more recent decades are in darker shades, so the lightest circles are for 1940–1950 and the darkest for 2000–2010. Panels (a) and (b) include all variation in the data, national and location-specific changes, and pool all decades from 1940 to 2010. Panels (c) and (d) control for the national growth rate in each decade, thereby isolating location-specific changes. Panels (e) and (f) show growth rates for each CZ over the entire 70-year period, 1940 to 2010. The top left corner of each plot reports the slope and R^2 statistic from unweighted and population-weighted regressions.

These figures show the extent of the variation and the strength of the relationships between housing prices and wages. Ten percent faster wage growth is associated with 6.2–8.5% faster rent growth and 10–15% faster home price growth (depending on weighting). Large shares of the variation in rents and home prices are associated with variation in wages, with R^2 statistics ranging from 35–63%. These relationships are even stronger when isolating the location-level component by looking at relative growth in these prices. This is consistent with spatial equilibrium forces ensuring that, other things equal, places with higher wages must have higher prices of non-tradables such as housing in order to maintain equilibrium.

These conclusions are highly robust, including to using different sample periods (e.g., excluding the Great Recession, the housing boom and bust cycle, or the post-war economic boom), different market definitions (e.g., counties or MSAs), and different measures of income and housing costs (e.g., total family income or gross rents). They are robust to controlling for population changes and dropping migrants. They hold across the distribution of wages and housing costs (e.g., at the 25th or 75th percentiles). We find similar results if we restrict the analysis to the modal 5-room home or if we use repeat-sale transaction level data from Zillow or FHFA in the periods in which they are available. Placebo-type tests suggest that the results are not driven by any spurious relation or omitted variable. For example, we find that factors that are largely determined at the national level, such as Social Security income (a federal program) or electricity demand (supplied competitively on a national grid) are not associated with location-level wage changes. Details of all robustness tests can be found in Appendix B.

5 Implications for the Riskiness of Owning versus Renting Housing

5.1 Interpreting the Results

The implications of the strong empirical regularities documented in Section 4 for the riskiness of owning versus renting depend on the underlying factors driving these relationships. The main issues are the extent to which the variation reflects risk as opposed to predictable heterogeneity and the extent to which the overall patterns are likely to generalize to other contexts and time periods. Understanding the underlying mechanisms is crucial for addressing both of these issues.

The large, persistent differences in wage growth across locations and time periods are suggestive that labor productivity risk is substantial, especially at the local level. This interpretation is consistent with a large literature on the local labor market effects of shocks to the demand for tradable goods (e.g., Barro and Sala-i Martin (1991), Blanchard and Katz (1992), Autor, Dorn and Hanson (2013), Giannone (2018), Howard and Liebersohn (2018)).[10]

Labor demand shocks affect housing prices both by affecting the cost of producing housing (since labor is an input to housing production) and, likely of greater quantitative importance, by shifting the demand for housing. The strength of the association between changes in wages and in housing prices suggests that wage changes shift the demand for housing substantially—likely due to some combination of a large income elasticity of demand for housing and spatial

[10]Additional support for the conclusion of substantial local labor demand risk is the strong positive correlation between the growth of a location's wages and rents relative to the national average, since shocks to local consumer amenities, for example, are a force toward rents covarying *negatively* with wages (Roback, 1982). The strong relationship between location-specific changes in wages and rents is suggestive of strong forces toward spatial equilibrium. The evidence from the 1940–2010 period suggests a fairly strong positive relationship between national-level changes in wages and rents as well, though this is necessarily less certain since only a single national time series is observed.

equilibrium-type responses to location-specific wage changes—and that the supply of housing is relatively inelastic. The inelasticity of housing supply in response to negative demand shocks is a natural consequence of the durability of housing (Glaeser and Gyourko, 2005). The inelasticity in response to positive demand shocks in the locations with the greatest wage growth over this period is a consequence of the increased difficulty of building in these locations between the 1960s and the 1990s, largely due to policies (Glaeser and Gyourko, 2018).

5.2 Riskiness of Owning versus Renting Housing

Section 4 showed that growth rates in the spot prices of labor and housing vary significantly across locations and over time. In this section, we aim to go beyond these summaries of the raw data to construct measures that are more closely connected to the key conceptual experiment introduced in Section 2: the ex ante risk in a household's lifetime net income from risk in market wages, rents, and home prices.

Our goal is to construct a transparent empirical counterpart of the key conceptual experiment introduced in Section 2. To do so, we imagine a household that at the beginning of adult life draws a present value of lifetime net income (income net of housing costs) from a known distribution. The household starts adult life at age 25, retires at age 65, and dies at age 95. The distribution of lifetime net income depends on whether the household owns or rents its home, given the different exposures of owners and renters to the risk in housing prices. Given (ex post) distributions of the household's lifetime net income by housing tenure status, we calculate the household's ex ante willingness to pay to eliminate the risk it faces separately by whether it is an owner or a renter.

Risk in the present value of lifetime net income comes from risk in the market wages, rents, and home prices during the household's life. We imagine a household that starts adulthood in a market with the median wage, median rent, and median home price across markets in 2010. The household then draws a vector of growth rates of wages, rents, and home prices over time from the observed empirical distribution of growth rates of these prices over time during the 1940–2010 period in each of the 721 CZs. For example, a renter household that drew the CZ that includes New York City for life would have present value lifetime housing costs of

$$R = \sum_{t=1940}^{2010} \frac{g_{1940,t}^{NYC} R_{2010}^{med}}{(1+r)^{t-1940}},$$

where $g_{1940,t}^{NYC} = R_{t,NYC}/R_{1940,NYC}$ is the gross growth rate in the median rent in the CZ that includes New York City from 1940 to t and R_{2010}^{med} is the median rent in the U.S. in

2010. Income during the working life equals the market wage in each period. Income during retirement equals 50% of the household's age-64 income. The probability of drawing each CZ is proportional to its population.

Homeowners face transactions costs of 3% of the home price when they buy or sell their home. We use three scenarios for dealing with unobserved carrying costs. The first two set carrying costs such that the expected present value of housing costs is the same for renters and owners. The implied ratio of carrying costs to home prices is 2.1%, which is slightly higher than the depreciation rate of the housing stock of about 1.8% typically used in calibrations in the literature.[11] To assess the risk aspect of carrying costs, we report two extreme cases. In case (a) we assume that carrying costs are 2.1% of national-level home prices but otherwise uncorrelated with the local housing market. In case (b) we assume that carrying costs are 2.1% of the local home price and therefore perfectly correlated with the local housing market. We also report a third case, case (c), in which we assume that carrying costs are zero.

Risk in the household's net income comes from risk in the rates at which wages, rents, and home prices in its location(s) grow relative to the national average. Ex post lifetime net income is greater in states of the world in which the household's location experiences faster wage growth, slower rent growth (if the household rents), or faster house price growth (if the household owns). This risk process isolates the across-location risk in the growth rates of wages, rents, and home prices over time; it excludes aggregate risk by conditioning on the single, observed realization of aggregate growth. It is a highly flexible model of the underyling data generating process.

[Table 2 about here.]

Table 2 reports statistics about the levels of, variation in, and correlations between the resulting lifetime present values. Panel A summarizes the lifetime present values of gross income and housing costs. There is substantial variation in these lifetime measures driven by the substantial, persistent differences across locations in the growth rates of the underlying prices. The standard deviation of earnings is $285,000 (coefficient of variation of 12%), of a renter's housing costs is $130,000 (coefficient of variation of 28%), and of an owner's housing costs is about $100,000.

Panel B summarizes the lifetime present values of the *net* income of renters, owners, and "pre-payers," i.e., households with full, actuarially fair housing cost insurance. Owners face the greatest net income risk, with a standard deviation of $322,000 and CV of 17%.

[11] This higher rate likely reflects other factors not accounted for by our hedonic regressions, such as unobserved quality differences between owner-occupied and rental units, and frictions in the rental market such as moral hazard.

Renters face the smallest net income risk, with a standard deviation of $209,000, $113,000 smaller than that of owners' and $76,000 smaller even than that of pre-payers with no housing exposure. That is, the renter's exposure to housing costs *reduces* its consumption risk (relative to having no exposure to housing cost risk). This is despite the fact that rent risk is substantial, as Sinai and Souleles (2005) emphasize. The renter's lifetime housing costs have a standard deviation of $130,000 and a coefficient of variation of 28%. Despite the substantial risk in a renter's housing costs, exposure to these costs reduces net income risk because of the strong correlation between rents and wages. The owner's exposure to housing costs, by contrast, increases its consumption risk substantially, independent of carrying costs. These different effects on consumption risk are the result of the different correlations of the lifetime housing exposures with lifetime wages. The strong positive correlations between wages and housing prices create a positive correlation between the earnings and housing costs of renters of 74% and a negative correlation between the earnings and housing costs of owners of −22% (see Appendix Table A.4).

These calculations show that renting reduces the ex post dispersion in lifetime net income relative to owning and even relative to having no housing exposure, given the substantial, correlated wage risk facing households. We now turn to quantifying the welfare impact. We measure the welfare cost of different risk exposures as the compensating variation version of the certainty equivalent, i.e., the household's ex ante willingness to pay to eliminate all of the risk it faces (in income as well as in housing costs),

$$V_0\left(E(Y) - E(H) - CV\right) = E\left[V_0\left(Y - H\right)\right],$$

where $V_0(x)$ is ex post indirect lifetime utility as a function of lifetime net income x, Y and H are the lifetime present values of gross income and housing costs, respectively, and CV is the compensating variation. We assume that lifetime utility is a standard, additively separable power function with exponential discounting: $V_0(x) = \sum_{t=1}^{T} \beta^{t-1} c_t^*(x_t)^{1-\sigma}/(1-\sigma)$, where x_t is per-period net income ($\sum_{t=1}^{T} \beta^{t-1} x_t = x$), $\beta = 1/1.03$ is the discount factor, and σ is the coefficient of relative risk aversion, which is three in our baseline calculations. In this static framework, the household learns its lifetime net income immediately and perfectly smooths consumption within each (lifetime) state of the world. This tends to reduce the welfare cost of risk relative to the more realistic case in which information about the state of the world is revealed only gradually over time.

[Table 3 about here.]

Table 3 shows the results. An owner is willing to pay significantly more to eliminate risk than a renter, $58,000 (3% of consumption) versus $28,000 (1.5% of consumption).

Unsurprisingly given that renting significantly reduces the standard deviation of lifetime net income, with this risk process exposure to rent risk is valuable. The renter is willing to pay significantly less to eliminate risk than a household with no housing cost risk exposure (i.e., with full, actuarially fair housing cost insurance) would be: $28,000 (1.5% of consumption) versus $47,000 (2.5% of consumption). The risk to lifetime net income from differential growth in market wages is substantial, so the strong correlation between differential growth in market wages and rents means that renters' short position with respect to housing provides a highly valuable hedge against their income risk from market wage risk.

These results suggest that the particular place and time in which a household lives has a large impact on its lifetime earnings and housing costs, and so on its living standard. Crucially, this "luck of the draw" is much more consequential for owners than renters for two key reasons. First, for a household that does not have a long horizon in its home, its lifetime housing costs tend to be more variable by owning than by renting. The much greater variation in home prices than rents is why the "direct effect" of housing costs on consumption risk is in many cases larger for owners, though, as emphasized in Section 2, a sufficiently long horizon in a home reduces an owner's present value sales price risk toward zero.

Second and more important, lifetime housing costs are strongly positively correlated with lifetime income for renters and weakly negatively correlated for owners. These correlations mean that renting hedges income risk and owning exacerbates it. As a result, for working households, a renter's exposure to rent risk is much less costly than it otherwise would be—and in fact appears to have negative cost, i.e., a positive value, for a typical household—and an owner's exposure to home price risk is much more costly than it otherwise would be. In other words, the "portfolio effect" terms introduced in Section 2 reduce the riskiness of renting and increase the riskiness of owning. This suggests that for the typical working household, owners' hedge against rent risk is a cost, not a benefit, of owning. Ignoring the strong correlations between housing costs and the rest of working households' portfolios therefore systematically understates the riskiness of owning and overstates the riskiness of renting.

Although these calculations have several attractive features—especially their transparency, close connection to the data, and the minimal structure they impose on the underlying data generating process—they abstract from many factors that could influence the welfare cost of owners' and renters' risk exposures. Most notably, they use a simple static model in which the household learns its full lifetime income and housing costs immediately at the beginning of the life cycle and faces only a present value budget constraint. As a result, the household can adapt its consumption to the realized state of the world perfectly; the risky realization of the present value of consumption is smoothed perfectly within each state of the world, without any additional costs from a gradual resolution of uncertainty or borrowing constraints. In reality, households learn about the realizations of market wages and housing prices only gradually over time. By assuming instead that households learn their net income immediately, these calculations tend to understate the welfare cost of these risks and the risk benefit of renting over owning, since the greater the risk, the greater the uncertainty about the state of the world not only from "behind the veil" but also during one's lifetime, and so the greater the difficulty of adapting one's choices to the state of the world. These calculations also ignore within-market, idiosyncratic risk in wages, rents, and home prices, which tends to understate the welfare cost of each of these risks in isolation, overstate the share of income risk that is hedged or exacerbated by different housing positions, and understate the net risk benefit of renting, since idiosyncratic risk in home prices is quite large, both absolutely and relative to plausible levels of idiosyncratic risk in rents. The calculations therefore likely understate the net risk benefit of renting.

For these reasons and others, we also quantify the welfare cost of owners' and renters' exposures in a rich life cycle model. The results, shown in Appendix Table A.3, are quite similar quantitatively and identical qualitatively: For a typical household owning is much riskier than renting. Details are in Appendix C.[12]

5.3 Heterogeneity in the Riskiness of Owning versus Renting

These results about the riskiness of owning versus renting for a "typical household," together with the discussion of underlying mechanisms in Section 5.1, suggest important heterogeneity by household and location characteristics. Probably the most fundamental determinant of the relative riskiness of owning versus renting is the underlying risk process. As emphasized in Section 5.1, this depends on the interaction between the elasticity of housing supply and the riskiness of labor demand. A household whose labor market is exposed to greater productivity risk and whose housing market has less elastic supply will tend to face wages and housing prices that are both more volatile and more highly positively correlated. Such a household's housing tenure choice is a crucial determinant of the overall riskiness of its living standard.

Prime examples of markets with volatile labor demand and inelastic housing supply are

[12] Although the life cycle model can account for a variety of factors that our preferred calculations in the main text do not, such as borrowing constraints and the gradual resolution of uncertainty, they require much stronger assumptions on other dimensions, such as on the underlying data generating process and household's beliefs. Of particular importance is the data generating process. Each of the wide range of processes we have considered appears to miss key features of the data, such as its skewness or thickness of tails. Given the crucial role of the risk process for our question of interest, the more flexible, non-parametric modeling of this process that is possible in the static lifetime calculations in the main text is our main reason for adopting these as our baseline.

booming coastal cities with constraints that limit housing supply and declining "industry towns," like many locations in the Rust Belt of the U.S., since the durability of housing means that housing supply responds inelastically to negative demand shocks (Glaeser and Gyourko 2005). The extent to which a resident of Detroit in 1950, when Detroit's wages and rents were among the highest in the U.S., shared in Detroit's dramatic decline over the second half of the 20th century was much greater for owners than renters. Similarly, the extent to which a resident of Seattle in 1980, when Seattle's wages and rents were among the lowest in the U.S., shared in Seattle's dramatic rise over the subsequent decades was also much greater for owners than renters. For residents of different locations today, who do not know if their location will go the way of Detroit in the second half of the 20th century or of Seattle in the decades following 1980, renting provides a valuable hedge against wage risk, whereas owning "doubles down" on wage risk.

In terms of household characteristics, the results suggest that a crucial determinant of the riskiness of owning versus renting is a household's exposure to wage risk. Two-earner households have significant exposure to wage risk and so much to gain from hedges against it. Retirees have relatively little exposure to wage risk, so for them the relative riskiness of owning versus renting depends on the trade-off emphasized by Sinai and Souleles (2005): Renters are exposed to rent risk and owners are exposed to sales-price risk. These considerations suggest that from a risk-management perspective, the optimal life cycle profile of housing tenure is likely renting during prime working years when exposure to wage risk is greatest, before owning beginning around the time of retirement (or, better yet from a risk perspective, limiting one's housing exposure by taking out an appropriate reverse mortgage).

Another important household characteristic is the relation between current and future housing demand. One determinant of this is the household's horizon in its home or location. As Sinai and Souleles (2005) emphasize, the longer the horizon, the greater the exposure to rent risk and, eventually, the smaller the present-value exposure to home price risk. But unlike in Sinai and Souleles (2005), in which rent risk is unambiguously bad and a longer horizon unambiguously increases the insurance advantage of owning, our analysis suggests that for a household whose income is sufficiently positively correlated with rents, a longer horizon does not necessarily increase the relative insurance advantage (reduce the insurance disadvantage) of owning, since it also increases the value of the rent risk hedge against wage risk.[13]

[13] Sinai and Souleles (2005) show that people tend to move to locations that have experienced price growth similar to that of the locations from which they came. If housing costs and wages in a location were uncorrelated with one another, such moves to correlated markets have a similar effect on the riskiness of owning versus renting as a lengthening of the horizon in a home. But because housing costs and wages are strongly positively correlated, for households that plan to work in the future, housing costs in locations a household might wish to live in the future are not a major source of risk, since greater housing costs tend to

Similarly, an owner of a small home who anticipates having a much greater demand for housing in the future may be not long but short housing. Compared to renting, owning tends to reduce his net (short) exposure to housing cost risk. The question for optimal portfolio management, of course, is how any given exposure contributes to the riskiness of the overall portfolio. For many prime-age working households, short positions with respect to housing likely reduce the riskiness of the overall portfolio.

6 Policy Implications

Our results have implications for two main sets of policies. The first is the many policies that restrict the supply of housing, especially in booming coastal locations (see, e.g., Glaeser and Gyourko, 2018). Such restrictions, by reducing the elasticity of housing supply, greatly increase the extent to which labor demand shocks capitalize into wages and housing costs and so the risk in earnings and housing costs. They also greatly raise the risk consequences of housing tenure choices, by increasing both the value of a renter's hedge against wage risk and the cost of an owner's exposure to home price risk. Our results suggest that the risk cost of such policies could be substantial. Better understanding these costs is an important topic for future research.

The second set of policies our results have implications for are the many policies that implicitly subsidize owner-occupied relative to rental housing. Although a full welfare analysis of these policies is beyond the scope of this paper, here we report simple calculations aimed at providing a rough sense of the likely magnitudes. The key idea is that any tendency of households to underestimate the risk (or other) costs of owning relative to renting magnifies the efficiency cost of policies that explicitly or implicitly subsidize owning relative to renting (above and beyond the level appropriate to internalize any externalities).

A variety of evidence suggests that many households underestimate the private costs of owning relative to renting. Those who teach economics are likely familiar with the widely-held view that renting is "throwing money away" in a way that owning is not, which reflects a failure to understand the opportunity cost of living in a home you own (the foregone opportunity to rent it to someone else). More germane to our findings, many people appear to underestimate the relative riskiness of owning. For example, recent survey evidence suggests that many people view housing as an extremely safe investment, safer even than Treasury bonds (Adelino et al., 2018). That owning is safer than renting is a staple of financial planning advice (e.g., Malkiel (1999)) and is consistent with the prevailing view in urban economics. In light of our finding that for many households owning is riskier than renting,

be compensated by greater wages and vice-versa.

the widespread view that owning is safer suggests that many people are underestimating the riskiness of owning relative to renting. To the extent that misperceptions about risk or other costs of owning affect people's choices, they tend to reduce welfare both directly, by leading people to make choices that are privately sub-optimal, and indirectly, by magnifying the efficiency cost of policies that subsidize owner-occupied relative to rental housing.

As a conservative benchmark, we ignore any bias from failures to appreciate the opportunity cost of owning and suppose that the only bias is that people underestimate the riskiness of owning relative to renting. We suppose—again, aiming to be conservative—that people understand that owning, like renting, is risky, but fail to account for the correlations between wages and housing costs. In that case, our results suggest that a "typical" household would underestimate the *lifetime* (risk) cost of owning relative to renting by about 1.6 percent of lifetime net income, or about \$16,000 for a household with \$1m of lifetime net income.[14]

A variety of policies subsidize owner-occupied relative to rental housing above and beyond what externalities appear to justify. The implicit net subsidy on owner-occupied housing (beyond that justified by externalities), σ_{sub}, depends on many particulars of a household's situation, including its marginal income tax rate, whether it itemizes deductions, and the extent to which it will be a net seller of housing before death. But the key policy that implicitly subsidizes owner occupied housing is that owner occupiers' implicit rent—the rent they implicitly pay their landlord self as their tenant self—is not taxed in most countries, whereas explicit rent from non-owner tenants to their landlords is.[15] A rough measure of the implicit subsidy on owner-occupied housing is the product of a typical household's marginal income tax rate and its (possibly implicit) rent, τR. As an approximation, the annual rental

[14]The relevant results in Appendix Table A.3 are that (1) the true risk advantage of renting, accounting for correlations between earnings and housing costs, is worth 1.1% of lifetime net income and (2) the perceived risk advantage of renting, i.e., the risk advantage if earnings and housing costs were uncorrelated, is worth −0.5% of lifetime net income (i.e., renting is perceived to be riskier). Combining these suggests a total "internality wedge" between actual and perceived risk costs of owning versus renting of 1.1+0.5=1.6% of lifetime net income.

[15]Many other policies also implicitly subsidize owner-occupied relative to rental housing, including implicit taxes on high-density housing (and so subsidies on low-density housing, which is more likely to be owner-occupied) such as zoning restrictions and government-provided infrastructure. Tax provisions related to housing are varied and often poorly understood. The "tax expenditure" associated with the lack of taxation of owner occupiers' implicit rent totaled \$117b in 2018 (https://www.treasury.gov/resource-center/tax-policy/Documents/Tax-Expenditures-FY2019.pdf), making it the second-largest tax expenditure after that for employer contributions for health insurance. The deductibility of mortgage interest and property taxes from a homeowner's taxable income do not as a first approximation subsidize owner-occupied housing relative to rental housing because landlords can make similar deductions. The other main difference is that landlords can deduct maintenance and depreciation, whereas owner occupiers cannot. Additional complications arise related to capital gains and the ability to defer these through purchases of other properties, etc. These many factors make it impossible to capture the full richness of the tax treatment of housing in a simple calculation, so our calculations aim only to be rough, back-of-the-envelope estimates of the likely magnitude for certain households.

yield on housing is about 10 percent (i.e., the market rent is about 10 percent of the capital price, Begley, Loewenstein and Willen (2019)), so the rent of a home worth the median home price of $220,000 in 2010 is about $22,000 per year. Taking a benchmark marginal tax rate on income of 25 percent, the implicit subsidy on owner-occupied housing is roughly $5,000 per year. As a rough measure of the implicit subsidy *per lifetime* of living in owner-occupied housing, the value of a 50-year annuity of $5,000 per year, using a 3 percent annual real interest rate, is $130,000.

As is well known, pre-existing distortions magnify the marginal excess burden of additional distortions that reinforce them; the marginal excess burden is not the usual triangle but the potentially much larger trapezoid. That many people appear to underestimate the cost of owner-occupied relative to rental housing magnifies the marginal excess burden of implicit housing subsidies. Similarly, the subsidies increase the welfare cost of misperceptions.

A simple measure of the marginal excess burden of the net subsidy on owner-occupied housing is

$$MEB_{sub} = \left(\sigma_{int} + \frac{\sigma_{sub}}{2}\right) \times \Delta Q_{sub},$$

where σ_{int} is the internality wedge (i.e., the subsidy that would increase the quantity of owner-occupied housing by as much as misperceptions do), σ_{sub} is the implicit net subsidy on owner-occupied relative to rental housing beyond that justified by externalities, and ΔQ_{sub} is the increase the quantity of owner-occupied housing from the implicit net subsidy. Table 4 shows the results.

While the absolute size of the efficiency cost depends on the quantity response, the implicit subsidies alone are sufficient to calculate the relative increase in the efficiency cost of the subsidy from misperceptions and of misperceptions from the subsidy. Misperceptions increase the efficiency cost of the subsidy by 25 percent, and the subsidy increases the welfare cost of misperceptions by a factor of 17.[16]

As a rough measure of the absolute size of the efficiency costs, assume as an approximation that for a typical household, the subsidy increases the fraction of its lifetime it owns as opposed to rents by 1/3, from the roughly 1/3 homeownership rate in Switzerland, which has small homeowner subsidies, to the roughly 2/3 homeownership rate in the U.S. (Assuming linear demand, this price responsiveness implies that misperceptions increase the fraction of time the typical household spends owning by about 4% of its lifetime.) In this case, the subsidy has a marginal excess burden of about $\$(16,000 + 130,000/2) \times 1/3 \approx \$27,000$

[16]Misperceptions increase the efficiency cost of the subsidy by $MEB_{sub}/MEB_{sub|\sigma_{int}=0} = 1+2\sigma_{int}/\sigma_{sub} \approx 1 + 2 \times \$16,000/\$130,000 = 1.25$). The subsidy increases the welfare cost of misperceptions by a factor of $1 + 2\sigma_{sub}/\sigma_{int} \approx 17$.

per household over its lifetime, whereas without misperceptions it would have been about \$21,700. Misperceptions themselves have a marginal efficiency cost of about $\$(130,000 + 16,000/2) \times 4\% \approx \$5,700$ per household over its lifetime.

These calculations, despite their many caveats, are suggestive that misperceptions about the costs of owning relative to renting have large costs, both direct, by leading people to make choices that are privately sub-optimal, and indirect, by magnifying the efficiency cost of policies that subsidize owner-occupied relative to rental housing. Additional evidence on people's perceptions of the costs of different housing positions and the extent to which these perceptions influence their choices is an important topic for future research.

7 Conclusion

A widely-held view expressed in the economics literature, in practical financial advice, and among the public at large is that owning one's home is safer than renting. In this paper, we show that for many households the opposite is true—renting is safer than owning, often significantly so. The key factor driving the difference relative to earlier work is the strong correlations between housing costs (rents and home prices) and wages at life-cycle-relevant frequencies. These strong correlations mean that for many working households, the exposure to rent risk that they would have as renters would provide a valuable hedge against their exposure to wage risk through their labor earnings. By contrast, the exposure to sale price risk that they would have as homeowners would increase the cost of their exposure to wage risk.

The analysis suggests many important dimensions of heterogeneity in the riskiness of owning versus renting. The key household characteristic is exposure to wage risk, mainly through labor supply. The key location characteristics are the volatility of labor demand and the elasticity of housing supply. A more detailed analysis of such heterogeneity is an important topic for future research.

These findings have important implications for economic research and policy. In terms of research, the strong correlations between wages and housing prices mean that analyses of earnings risk should account for housing cost risk and vice versa. Failing to account for these linkages can lead to misleading conclusions about the costs of different risk exposures and the fit of different models of household behavior. In terms of policy, our results, together with evidence that many households underestimate the relative riskiness of owning, suggest that the efficiency cost of many housing-related policies are larger than previously appreciated. Our results also suggest that the many policies that restrict the supply of housing could have large costs in terms of increasing risk in earnings and housing costs. Better understanding these issues is a high priority for future research.

References

Adelino, Manuel, Antoinette Schoar, and Felipe Severino, "Perception of House Price Risk and Homeownership," *NBER Working Paper No. 25090*, 2018.

Ameriks, John, Joseph Briggs, Andrew Caplin, Matthew D Shapiro, and Christopher Tonetti, "Long-Term-Care Utility and Late-in-Life Saving," *NBER Working Paper No. 20973*, 2015.

Autor, David H., David Dorn, and Gordon H. Hanson, "The China Syndrome: Local Labor Market Effects of Import Competition in the United States," *American Economic Review*, 2013, *103* (6), 2121–68.

Barro, Robert J. and Xavier Sala i Martin, "Convergence Across States and Regions," *Brookings Papers on Economic Activity*, 1991, *1*, 107–182.

Begley, Jaclene, Lara Loewenstein, and Paul S. Willen, "The Price-Rent Ratio During the Boom and Bust: Measurement and Implications," *Working Paper*, 2019.

Blanchard, Olivier J. and Lawrence F. Katz, "Regional Evolutions," *Brookings Papers on Economic Activity*, 1992, *1*, 1–61.

Campbell, Sean D., Morris A. Davis, Joshua Gallin, and Robert F. Martin, "What moves housing markets: A variance decomposition of the rent–price ratio," *Journal of Urban Economics*, 2009, *66* (2), 90–102.

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

Davis, Morris A. and Erwan Quintin, "On the Nature of Self-Assessed House Prices," *Real Estate Economics*, 2017, *45* (3), 628–649.

_ and Stijn Van Nieuwerburgh, "Housing, Finance, and the Macroeconomy," in "Handbook of Regional and Urban Economics," Vol. 5, Elsevier, 2015, pp. 753–811.

Fella, Giulio, "A generalized endogenous grid method for non-smooth and non-concave problems," *Review of Economic Dynamics*, 2014, *17* (2), 329–344.

Giannone, Elisa, "Skilled-Biased Technical Change and Regional Convergence," *Working Paper*, 2018.

Glaeser, Edward L. and Joseph Gyourko, "Urban Decline and Durable Housing," *Journal of Political Economy*, 2005, *113* (2), 345–375.

_ and _ , "The economic implications of housing supply," *Journal of Economic Perspectives*, 2018, *32* (1), 3–30.

Howard, Greg and Jack Liebersohn, "What Explains U.S. House Prices?Regional Income Divergence," *Working Paper*, 2018.

Hurst, Erik, Benjamin J. Keys, Amit Seru, and Joseph Vavra, "Regional Redistribution through the US Mortgage Market," *American Economic Review*, 2016, *106* (10), 2982–3028.

Ivanov, Ventzislav and Lutz Kilian, "A Practitioner's Guide to Lag Order Selection for VAR Impulse Response Analysis," *Studies in Nonlinear Dynamics & Econometrics*, 2005, *9* (1).

Jorgenson, Dale and Robert E. Hall, "Tax policy and Investment Behavior," *American Economic Review*, 1967, *57* (3).

Karahan, Fatih, Fatih Guvenen, Serdar Ozkan, and Jae Song, "What Do Data on Millions of US Workers Reveal About Life-Cycle Earnings Risk?," *Working Paper*, 2015.

Lockwood, Lee M., "Incidental Bequests and the Choice to Self-Insure Late-Life Risks," *American Economic Review*, September 2018, *108* (9), 2513–2550.

Malkiel, Burton Gordon, *A Random Walk Down Wall Street: Including a Life-cycle Guide to Personal Investing*, WW Norton & Company, 1999.

Nakajima, Makoto and Irina A. Telyukova, "Reverse Mortgage Loans: A Quantitative Analysis," *Journal of Finance*, 2017, *72* (2), 911–950.

Plazzi, Alberto, Walter Torous, and Rossen Valkanov, "Expected Returns and Expected Growth in Rents of Commercial Real Estate," *Review of Financial Studies*, 2010, *23* (9), 3469–3519.

Roback, Jennifer, "Wages, Rents, and the Quality of Life," *Journal of Political Economy*, 1982, *90* (6), 1257–1278.

Ruggles, Steven, Sarah Flood, Ronald Goeken, Josiah Grover, Erin Meyer, Jose Pacas, and Matthew Sobek, *IPUMS USA: Version 8.0 [dataset]*, Minneapolis, MN: IPUMS, 2018.

Sinai, Todd and Nicholas S. Souleles, "Owner-Occupied Housing as a Hedge Against Rent Risk," NBER Working Papers 9462, National Bureau of Economic Research, Inc January 2003.

_ and _ , "Owner-Occupied Housing as a Hedge Against Rent Risk," *Quarterly Journal of Economics*, 2005, *120* (2).

Figure 1 – Correlation of Wages, Rents and Home Prices

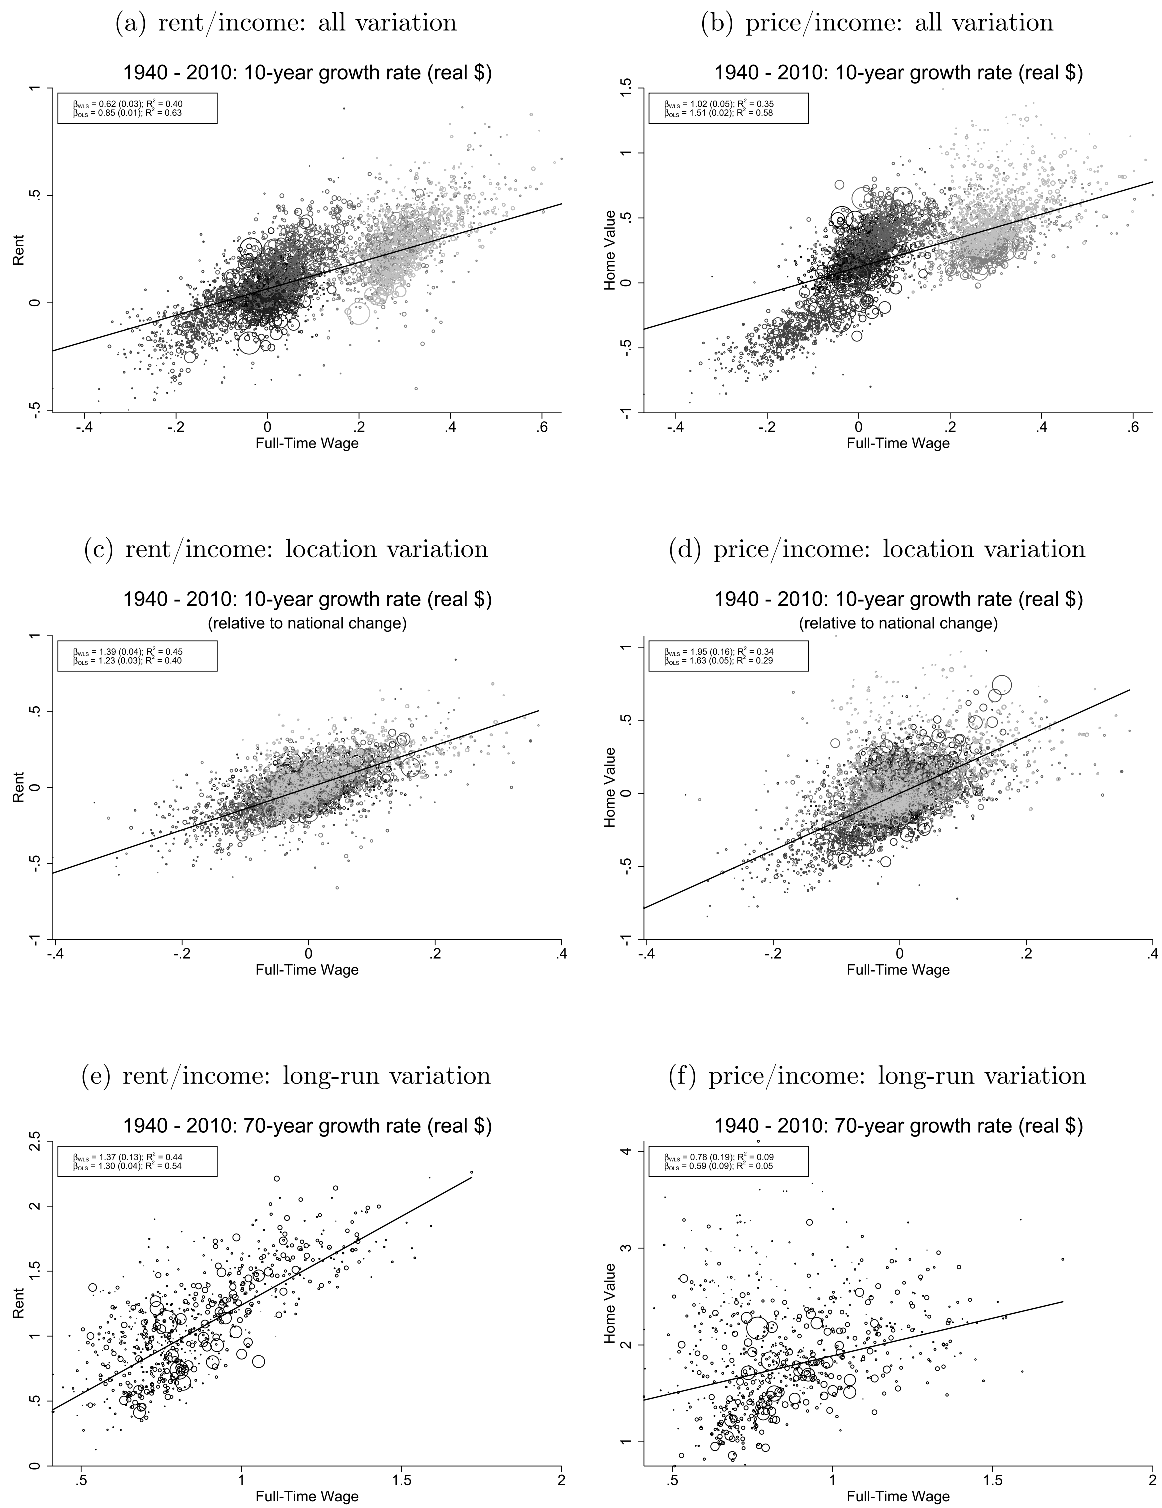

Notes. Log changes of median full-time wages, rents and home values by Commuting Zone (CZ). Circle areas represent CZ population size relative to total population in each decade. Observations in darker shades are from more recent decades. Source: Decennial Census and ACS.

Figure 2 – Income Hedge of Housing Tenure Choice

(a) rent/income: location variation

(b) price/income: location variation

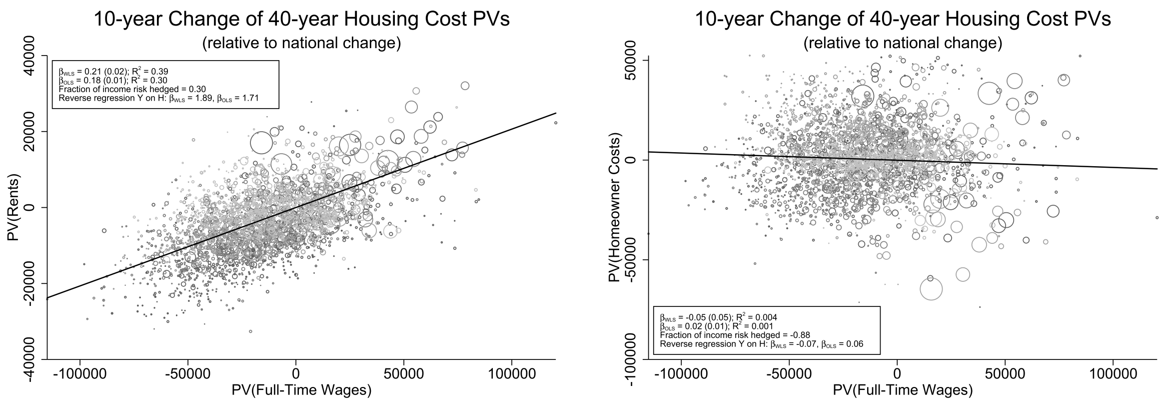

Notes. 10-year changes of 40-year present values (PVs) of full-time wages, rents and homeowner costs using a 3% discount rate, relative to the national average changes using Decennial Census and ACS data. Homeowner costs are the present value of carrying costs (see Appendix A.1.3) and realized capital losses. Circle areas represent CZ population size relative to total population in each decade. Observations in darker shades are from more recent decades.

Table 1 – Summary statistics of wages, rents and home prices

	Obs	Mean	Median	StDev	CV	p10	p90	Corr(x,wage)
	A. Levels							
2010 cross-section								
full-time wage income	721	51,313	50,327	8,913	17.4%	41,263	61,894	100.0%
rent	721	9,318	8,772	2,946	31.6%	6,000	13,836	80.7%
home value	721	235,631	187,500	129,090	54.8%	112,500	450,000	80.4%
	B. Log changes							
10-year log changes								
full-time wage income	5047	0.12	0.07	0.16	132%	-0.06	0.32	100%
rent	5047	0.14	0.13	0.16	113%	-0.06	0.33	61.7%
home value	5047	0.24	0.29	0.29	118%	-0.13	0.55	56.9%
10-year log changes, location-level only								
full-time wage income	5047	0	0.00	0.06	--	-0.07	0.08	100%
rent	5047	0	-0.01	0.13	--	-0.16	0.16	64.0%
home value	5047	0	-0.03	0.22	--	-0.24	0.29	52.0%
40-year log changes								
full-time wage income	2884	0.44	0.42	0.36	83.3%	-0.04	0.90	100%
rent	2884	0.60	0.57	0.32	53.4%	0.21	1.02	76.0%
home value	2884	0.93	0.82	0.52	55.8%	0.33	1.63	74.4%
40-year log changes, location-level only								
full-time wage income	2884	0	-0.02	0.12	--	-0.13	0.16	100%
rent	2884	0	-0.01	0.25	--	-0.30	0.31	69.4%
home value	2884	0	-0.03	0.36	--	-0.43	0.46	47.5%
70-year log changes, location-level only								
rent	721	0	-0.05	0.36	--	-0.42	0.50	69.7%
home value	721	0	0.05	0.47	--	-0.59	0.56	35.9%

Notes: An observation is a commuting-zone year (CZ-year) and all calculations use population weights. Values are in real dollars of 2010 after making top and bottom codes consistent over time.

Table 2 – Lifetime present values

	Obs	Mean	Median	StDev	CV	p10	p90
	A. Wages and Housing Costs						
full-time wage income, Y	721	2,400,884	2,334,108	284,557	11.9%	2,164,393	2,730,747
rent, R	721	465,653	435,342	128,113	27.5%	343,015	652,533
homeowner cost, $-P_T^{net}$							
a. uncorrelated carrying costs	721	465,653	477,120	100,728	21.6%	368,345	553,615
b. carrying costs correlated with P_t	721	465,653	434,628	103,890	22.3%	383,329	570,020
c. no carrying costs	721	61,471	72,938	100,728	163.9%	-35,836	149,434
$P_{t=0}$	721	228,100	228,100	0	0%	228,100	228,100
$PV(P_{T=70})$	721	178,836	167,015	103,843	58.1%	88,154	279,153
	B. Consumption (net income)						
renter	721	1,935,231	1,915,580	208,920	10.8%	1,722,481	2,147,594
prepayer	721	1,935,231	1,868,456	284,557	14.7%	1,698,741	2,265,095
owner with							
a. uncorrelated carrying costs	721	1,935,231	1,857,092	321,970	16.6%	1,665,522	2,289,725
b. carrying costs correlated with P_t	721	1,935,231	1,897,000	255,312	13.2%	1,656,066	2,221,200
c. no carrying costs	721	2,339,412	2,261,273	321,970	13.8%	2,069,703	2,693,906

Notes: An observation is a commuting-zone (CZ) and all calculations use population weights. Values are in real dollars of 2010. Discount rate is 3%. Closing costs are 3% (i.e., round trip costs of 6%). The initial starting values of (Y,R,P) are the population-weighted averages of the 2010 cross section. Carrying costs are calculated so that expected housing costs are equal for owners and renters. The ratio of implied carrying costs to home value is 2.1%, which closely matches empirical estimates of the rate of depreciation of the stock of housing structures (Davis and Van Nieuwerburgh, 2015). It also matches independent estimates of expenditure shares of homeowners for maintenance, repairs, insurance, property taxes, and other expenses (Bureau of Labor Statistics, Table 51. Housing tenure and type of area: Shares of average annual expenditures and sources of income, Consumer Expenditure Survey).

Table 3 – Willingness to pay to avoid consumption risk

	A. Relative risk aversion (σ)						
	0.5	1	2	3	4	5	6
Renter	5,204	10,177	19,542	28,292	36,584	44,544	52,277
Prepayer	9,276	17,918	33,583	47,484	60,003	71,437	82,013
Owner with							
a. uncorrelated carrying costs	11,670	22,421	41,578	58,191	72,831	85,935	97,839
b. carrying costs correlated with P_t	7,967	15,970	40,039	284,187	1,227,638	1,602,866	1,726,254
c. no carrying costs	9,832	18,980	35,499	50,036	62,985	74,654	85,290
	B. As a fraction of consumption						
Renter	0.3%	0.5%	1.0%	1.5%	1.9%	2.3%	2.7%
Prepayer	0.5%	0.9%	1.7%	2.5%	3.1%	3.7%	4.2%
Owner with							
a. uncorrelated carrying costs	0.6%	1.2%	2.1%	3.0%	3.8%	4.4%	5.1%
b. carrying costs correlated with P_t	0.4%	0.8%	2.1%	14.7%	63.4%	82.8%	89.2%
c. no carrying costs	0.4%	0.8%	1.5%	2.1%	2.7%	3.2%	3.6%
	C. Moving randomly every 10 years						
Renter	0.2%	0.3%	0.6%	0.8%	1.1%	1.3%	1.5%
Prepayer	0.2%	0.4%	0.7%	1.1%	1.4%	1.6%	1.9%
Owner with							
a. uncorrelated carrying costs	0.6%	1.1%	2.3%	3.4%	4.5%	5.7%	6.9%
b. carrying costs correlated with P_t	0.6%	1.1%	2.4%	3.9%	6.0%	9.5%	14.8%
c. no carrying costs	0.7%	1.4%	3.0%	4.9%	7.6%	11.9%	18.6%

Table 4 – Efficiency costs of subsidies and misperceptions

Panel A: Ingredients

Implicit subsidies ($s per lifetime)		Quantity response (in lifetimes)		
			β_L	β_H
Misperceptions, σ_{int}	16,000			
Net subsidy, σ_{sub}	130,000	$\Delta Q_{sub} = \beta\sigma_{sub}$	0.17	0.33
		$\Delta Q_{int} = \beta\sigma_{int}$	0.02	0.04

Panel B: Efficiency cost of subsidies ($s per lifetime)

	With misperceptions		Without misperceptions		Effect of misperceptions	
Formula	β_L	β_H	β_L	β_H	β_L	β_H
$\left(\sigma_{int} + \frac{\sigma_{sub}}{2}\right) \Delta Q_{sub}$	13,500	27,000	10,833	21,667	2,667	5,333

Panel C: Efficiency cost of misperceptions ($s per lifetime)

	With subsidies		Without subsidies		Effect of subsidies	
Formula	β_L	β_H	β_L	β_H	β_L	β_H
$\left(\sigma_{sub} + \frac{\sigma_{int}}{2}\right) \Delta Q_{int}$	2,831	5,662	164	328	2,667	5,333

Notes: σ_{int} is the internality wedge, i.e., the subsidy that would increase the quantity of owner-occupied housing by as much as misperceptions would under the assumptions discussed in Section 6. σ_{sub} is the implicit net subsidy on owner-occupied relative to rental housing beyond that justified by externalities. β_L (β_H) is the slope of demand for owner-occupied housing with respective to the subsidy rate assuming a low (high) price responsiveness. ΔQ_{int} (ΔQ_{sub}) is the increase in owner-occupied housing as a result of misperceptions (subsidies).

Online Appendix

The Riskiness of Owning Versus Renting Housing

Lee M. Lockwood   Scott R. Baker   Lorenz Kueng   Pinchuan Ong

A Appendix: Data

A.1 Location-Level Data and Sample Selection

We use the IPUMS 1940 100% full count data, the 1960 5% sample, the 1970 1% "Metro Form 1" and 1% "Metro Form 2" samples, the 1980 5% "State" sample, the 1990 5% "State" sample, the 2000 5% sample, and the 2011 5-year ACS 5% sample, all downloaded from https://usa.ipums.org/usa/data.shtml. The 1950 1% and the 2010 10% sample lack data on rent, income, and home values. We therefore use the 2011 5-year ACS for 2010 and we interpolate values in 1950 log-linearly between 1940 and 1960. When calibrating the model in Section C, we instead use a VAR-based imputation to appropriately reflect uncertainty between 1940 and 1960. Our results are robust to dropping 1950 from our sample or to using the shorter period from 1960-2010.

A.1.1 Defining Local Labor Markets for 1940-2010

The most detailed geolocation available in the microdata is a Public Use Microdata Area (PUMA), which maps either to a single county or a county group, and this assignment changes over time. To maintain confidentiality, only counties with a population larger than 100,000 (or 250,000 in 1970) are identified. The remaining counties are aggregated to county groups until these groups exceed this population threshold. We use separate allocation factors for each Decennial Census provided by IPUMS to allocate households to counties in cases where only county groups are identified. For example, if a PUMA consists of two counties, A and B, with population fractions 10% and 90% respectively, we allocate households in that PUMA to both counties, and multiply their household sampling weights by 10% in county A and 90% in county B. We then assign each county to a unique CZ using the definition of CZs in 1990.

A.1.2 Variable Construction

Home values are self-reported and available for owners only. Estimates include the full value of both house and land, even if the respondent resides on only part of the property. Condos were excluded in 1970 but included in 1980, 1990, 2000 and 2010. With the exception of 2010, the home value is not continuous and is reported as the midpoint of an interval. When measuring rent, we focus on contract rent. Gross rent includes utilities whereas contract rent may or may not depending on the household. Figure A.4 shows that utilities are not or only weakly correlated with local income changes. However, Figure A.5 shows that our results are robust to using gross rent instead, but gross rent is available only from 1960 on.

A.1.3 Homeowner's Carrying Costs

We use two extreme approaches to assess how incorporating carrying costs could affect the riskiness of owning. Both approaches set user costs such that expected costs from renting and owning are equalized on average, i.e., $E(R - P_T^{net}) = 0$. The first approach assumes that user cost changes are uncorrelated with house price changes, while the second assumes that they are perfectly correlated.

Approach 1 We start from the no-arbitrage condition between identical rental and owner-occupied properties over holding period T,

$$E_c\left[P_{c0} - \frac{P_{cT}}{(1+r)^T} + \sum_{t=1}^{T} \frac{O_{ct}}{(1+r)^t} - \sum_{t=1}^{T} \frac{R_{ct}}{(1+r)^t}\right] = 0, \tag{A.1}$$

where E_c are expectations in period 0 and location c over horizon T, which we estimate in the data by their unconditional mean in each location. In the Decennial Census, we do not observe identical rental and owner-occupied units. The average owner-occupied unit has a higher value than the average rental unit. We use the implied carrying costs to make them comparable,

$$EPV_c[O] \equiv E_c\left[\sum_{t=1}^{T} \frac{R_{ct}}{(1+r)^t} - \left(P_{c0} - \frac{P_{cT}}{(1+r)^T}\right)\right]. \tag{A.2}$$

We define owner's housing costs as

$$H_{own} = P_{c0} - \frac{P_{cT}}{(1+r)^T} + EPV_c[O]. \tag{A.3}$$

Approach 2 Starting from the 1-period no-arbitrage condition,

$$R_{ct} = P_{ct} - \frac{P_{c,t+1}}{(1+r)} + O_{ct}, \tag{A.4}$$

we define

$$E_c\left[\tfrac{O}{P}\right] \equiv E_c\left[\frac{R_{ct}}{P_{ct}} - \left(1 - \frac{P_{c,t+1}/P_{ct}}{1+r}\right)\right] \tag{A.5}$$

and owner's housing costs as

$$H_{own} = P_{c0} - \frac{P_{cT}}{(1+r)^T} + \sum_{t=1}^{T} \frac{P_{ct} \cdot E_c\left[\frac{O}{P}\right]}{(1+r)^t}. \tag{A.6}$$

A.2 Cash Flow versus User Cost Approach

The user cost of capital approach (Jorgenson and Hall 1967) equates the rental rate with the opportunity cost of owning the capital for one period,

$$R_{t+1} = \left[r + \tau^{prop} + \delta - \frac{\Delta P_{t+1}}{P_t}\right] P_t \tag{A.7}$$

$$\equiv (1 + r + o)P_t - P_{t+1}, \tag{A.8}$$

with rate of carrying cost $o = \tau^{prop} + \delta$. Solving this difference equation forward, we obtain

$$P_0 = \sum_{t=1}^{T} \frac{R_t}{(1+r+o)^t} + \frac{P_T}{(1+r+o)^T}. \tag{A.9}$$

Hence, the "user cost of capital approach" adjusts the denominator – the effective discount rate $r+o$ – while the "cash flow approach" adjusts the numerator, $R - O$,

$$P_0 = \sum_{t=1}^{T} \frac{R_t - O_t}{(1+r)^t} + \frac{P_T}{(1+r)^T}. \tag{A.10}$$

O are total property taxes and maintenance costs per period. For example, if tax values of houses are reassessed every period, then $O_t = (\tau^{prop} + \delta)P_t$.

B Appendix: Wage, Rent, and Home Price Variation

This section provides extensive robustness checks of our main findings in Section 4 and provides details about the data generating process we estimate to compute the individual welfare results in the life-cycle model of Section C.

B.1 Robustness of Income and Housing Cost Correlations

B.1.1 Placebo-type Tests using Social Security and Utility Costs

A reasonable concern is that the results might be driven by some spurious relation or omitted variable. To show that this is not the case, we redo the analysis using values that are largely determined at the national level, such as Social Security income, a federal program, or electricity demand, which is serviced competitively on a national grid, as well as all other components of utility costs. Figure A.4 shows that location-level income changes indeed have a negligible impact on these values.

B.1.2 Sample Periods

In Figure A.6 we drop the 2000s to test whether these correlations are driven by the recent housing boom and bust cycle and the Great Recession. We find that excluding this decade does not affect the correlations much. As mentioned in Section 3, the 1950 Decennial Census lacks information on housing costs. In the main analysis, we log-linearly interpolate values between 1940 and 1960, effectively giving this period the weight of two decades. Figure A.7 shows that dropping the 1940s after linear interpolation and thereby counting these two decades only once does not change our findings.

Similarly, Figure A.8 shows that these correlations are also robust to dropping the 1940s and 1950s and hence to dropping much of the post-war economic boom.

B.1.3 Definition of Local Labor and Housing Markets

Our results are robust to using alternative geographic definitions of labor and housing markets. To show this, Figure A.9 plots the same correlations as in Figure 1 but using counties instead of

CZs as locations. The strength of the correlation between wage income and housing costs is similar at the county and the CZ level.

B.1.4 Migration

One concern is that these correlations might be the outcome of economically selected migration and hence a reflection of increased cross-location economic segregation. We find that the correlations are largely unchanged when we control for population changes (Figure A.10) or drop migrants from the sample (Figure A.11).[17]

B.1.5 Housing Market Segmentation and Differences Across the Distribution

We also show that these correlations apply across the entire distribution of incomes and housing costs. Figures A.12 and A.13 show that the 25th percentiles and the 75th percentiles of the income and housing cost distribution move similar to the median, and Figure A.14 shows similar results for the mean (after making top and bottom codes consistent over time).

Figure A.15 tests whether our results are robust to potential housing market segmentation. For instance, it is conceivable that prices of different types of homes—say two and four room housing units—are set independently from each other, and the hedonic regression used in our main analysis might not sufficiently account for such market segmentation. To assess whether our findings are sensitive to such frictions, in Figure A.15 we limit the sample of housing units to those that have 5 rooms, which is the modal house size in all years. The figure shows that the correlations between income and housing costs are very similar as when we use the full sample of housing units.

B.1.6 Survey Responses versus Housing Transaction Data

Figure A.16 addresses the concern that there are systematic biases over time in survey responses about home values. Panels (a)-(c) show that Census home values line up very well against the transaction-based house prices from Zillow, both in dollars and in logs, as well as the county-level home price indices from the Federal Housing Finance Agency (FHFA), even better as the two transaction-based measures with each other as seen in panel (d). Similarly, panels (e) and (f) show that *changes* in median Census home values correlated as well with transaction-level price changes as do the FHFA and Zillow index with each other. This finding that self-assessed home values compare well with market prices is in line with recent results reported by Davis and Quintin (2017).

B.1.7 Family Income instead of Full-Time Salary

Figure A.17 extends our analysis to total family income, which reflects both changes in the market price of labor supplied as well as (changes in hours and labor market participation), financial income and government transfers. Total family income is as correlated with housing costs as is full-time salary. Hence, housing tenure choice provides households with an important opportunity to manage income risk that is not already covered by other financial assets in the household's portfolio or by the government (e.g., by the social insurance such as unemployment insurance, old-age pensions, or disability insurance).

[17] Migrants are households that report having moved to the county within 3 years prior to the survey (1 year in the 2011 5-year ACS).

B.1.8 Not Adjusting for Age and Size

Figure A.18 shows that not adjusting the for house or household age and size has very little effect on the correlations. While the hedonic regression substantially reduces the cross-sectional dispersion of full-time wages, home prices and rents, especially within locations, it as very little effect on their *growth rates*.

B.2 Extensions of Empirical Analysis

B.2.1 Riskiness of Owning vs Renting due to National-Level Changes

In Section 4 we focus on location-level risk by controlling for national-level changes with time fixed effects. Figure A.19 shows that we obtain similar benefits of renting to hedge earnings risk (respectively costs from amplifying earnings risk by owning) as in Figure 2 when we focus on national-level shocks instead of just location-level shocks.At the national level, we find quantitatively similar results. A \$100 decrease in the PV of life-time earnings is associated with an decrease in the PV of life-time rental expenditures of \$21, but with an *increase* in homeowner's costs of \$14. Similarly, the coefficients of the reverse regression are $\widehat{\beta}_{Y|R} = 4.4 \gg 1/2$ and $\widehat{\beta}_{Y|-P_T^{net}} = -0.41 \ll 1/2$, suggesting renting reduces total exposure to national-level earnings risk while owning increases it, consistent with the results using location-level risk. Stated differently, the fraction of national-level earnings risk that can be hedged by renting housing is 34% (compared to 30% for location-level earnings risk) while owning amplifies this risk by 61% (relative to 88% for location-level earnings risk); see equation **??**.

B.2.2 Correlated Carrying Costs

Figure A.20 shows that if we make the extreme assumption that changes in homeowner's carrying costs are perfectly correlated with changes in home prices and hence with changes in earnings, then these carrying costs provide some of the same insurance against earnings risk as does renting. This can be seen by the fact that the relation between location-level income changes and homeowner costs increases from -0.05 to 0.28 when using population weights and from 0.02 to 0.16 without weighting, and from -0.14 to 0.01 for national-level changes.

However, this increase is not enough to decrease the overall riskiness of owning as measured by equation **??**. The fact that home prices are much more volatile than rents largely outweighs the portfolio hedging effect of correlated carrying costs. While the amplification of national-level earnings risk due to owning housing decreases from 61% to 28%, the amplification of location-level earnings risk increases from 88% to 98%.

B.3 Data Generating Process (DGP) for Life-Cycle Model

We specify the following structural form for the DGP of annual wages, rents and home prices (in logs) for the life-cycle model in Section C:

$$y_{ct} = \eta_t + \eta_{ct} \qquad \text{(log of wage, rent and home price)} \tag{B.1}$$
$$\eta_t = A(t) + B_n(L)\eta_{t-1} + \epsilon_t^n \qquad \text{(national-level risk)} \tag{B.2}$$
$$\eta_{ct} = B_l \eta_{c,t-1} + \epsilon_{ct}^l \qquad \text{(location-level risk)} \tag{B.3}$$

The deterministic time effect $A(t)$ and the lag polynomial $B_n(L) = \sum_{i=1}^{p} B_{n,i} L^{i-1}$ can model different assumptions about the information/foresight that consumers have about macroeconomic trends. One extreme assumption is that $A(t) = 0$ and $B_n(L) = I_3$ so that national growth is unpredictable, i.e., follows a random walk. Another extreme assumption is $A(t) = A_t$, $B_n = V(\epsilon_t^n) = 0$ so that consumers know all future national-level changes, which are captured by the time fixed effects A_t, and hence there is no national-level uncertainty. An intermediate case is letting $B_n(L)$ be arbitrary and $A(t)$ being a polynomial (say a constant, or a linear or quadratic trend) such that trend growth is known (A) and also how fast the economy reverts back to trend (B_n) after a national-level shock ϵ^n.

Consumers estimate a reduced-form panel vector auto-regression (VAR) at 10-year frequency and do not differentiate between national- and location-level shocks. Instead, they estimate a first-order VAR on the combined process, which we show below is a good approximation to (B.1),

$$y_{ct} = \mu + \rho y_{c,t-1} + u_{ct}. \tag{B.4}$$

This greatly reduces the state space and increases computational speed when solving the life-cycle model. More specifically, having location-level data from the Decennial Census every 10 years from 1940-2010 and annual data from the BLS from 1964-2017 shown in Figure A.3, we proceed as follows:

1. We estimate B_l and $V(\epsilon_{ct}^l)$ in (B.3) using Decennial Census data and a panel VAR with time fixed effects to control for the national-level process (B.2). We find that the estimated B_l is very close to the identity matrix and we therefore impose it for simplicity.

2. Based on Schwarz's Bayesian information criterion (SBIC) applied to the annual BLS data,[18] we estimate a second-order VAR in (B.2).[19]

3. We simulate 1,000 years of data from this second-order VAR in (B.2) using 1,000 log-normal draws with estimated covariance matrix $V(\epsilon_t^n)$. We then keep every 10th observation. Applying the SBIC again suggests that a first-order VAR is the best representation of this decadal process, which justifies the reduced-form representation in (B.4).

4. We then simulate a panel of 500 locations, each with 80-years of data, based on (B.1)-(B.3). Each location starts from the same initial level of income, rent and home price, which we set to values roughly equal to 1960 values.[20] Specifically, we start the process with household income of $25,000 in 1960. We set rents equal to $5,000 corresponding a median rent-to-income ratio of 0.2, the average in our 1940-2010 Census sample. We set initial home value to $90,000 such that the price-income ratio is 3.6, which again matches the historical average in

[18]All information criteria suggest a VAR of order 2 if we use full-time earnings (i.e., median usual weekly earnings of full-time workers 16 years and older), which is only available from 1979 on. When we use hourly wage (i.e., average hourly earnings of production and non-supervisory employees), which goes back to 1964, then SBIC suggests a second-order VAR while AIC and HQIC suggest very long lag orders of about 10. Ivanov and Kilian (2005) find that for sample sizes smaller than 120, as in our case, the SBIC is the most accurate, hence we choose a lag order of 2.

[19]The constant is small and never statistically significant, hence we drop it when simulating data below. We also tested for co-integration. With an underlying VAR of order 2, we cannot reject that the process is not co-integrated. If we instead impose that the underlying annual process is a VAR of order 1, then we find evidence of one co-integrating equation. Since this evidence is weak, we therefore opt for a second-order VAR instead of a vector error correction model (VECM).

[20]We choose 1960 because this is the midpoint of our location-level Decennial Census sample of 40-year PVs used in Figure 2, which contains PVs from 1940-1970, 1950-1980, 1960-1990, 1970-2000, and 1980-2010. By choosing the midpoint of this sample, we do not have to take a strong stand on expected real growth in the economy. The dollar values of the national-level and location-level 40-year PVs in Figure 2 are therefore comparable.

the Census sample. This home value implies a price-rent ratio of 18.[21] The log of these three numbers is the starting value η_0 and $\eta_{c,0} = 0$ of all locations. We then estimate the panel VAR representation (B.4) of this process, which yields estimates of μ, ρ and $V(u_{ct})$. This is the main specification of the DGP for the life-cycle model in Section C.

C Appendix: Riskiness of Owning Versus Renting in a Dynamic Life Cycle Model

Section 5.2 computes the welfare impact of correlations between risks in a simple static model. In particular, the model assumes that households learn their full lifetime income and housing costs at the beginning of the life cycle and face only a present value budget constraint. Here, we consider a dynamic model in which information about the state of the world is revealed only gradually over time. We follow as closely as possible the key modeling assumptions of related work (e.g., Davis and Van Nieuwerburgh 2015, Hurst, Keys, Seru and Vavra 2016, Nakajima and Telyukova 2017) and test the robustness of the conclusions to plausible alternatives. Results obtained are similar to our main results.

C.1 Model

There are two types of individuals in this model: renters and owners. Individuals live at most eight periods, which correspond to decades (age decades 20s to 90s), with non-zero probability of dying every period.[22] Each period, individuals might be hit by a "tenure shock" that shifts renters into ownership and vice versa. After the tenure shock materializes, individuals choose how much to consume and save. Individuals have access to a consumption floor guaranteed by fiscal transfers. Upon death—which happens with certainty by age 100—the agent leaves a bequest, which includes the sale value of the house for owners. Expected utility in period t is[23]

$$EU_t = u(C_t) + \sum_{j=t+1}^{T+1} \beta^{j-t} \left(\prod_{s=t+1}^{j-1} (1-\delta_s) \right) \left[(1-\delta_j) E_t u(C_j) + \delta_j E_t v(B_j) \right], \tag{C.1}$$

where C_t is consumption, B_t is bequests, and δ_t is the exogenous probability of dying just before period t (so $(1-\delta_t)$ is the probability of surviving from period $t-1$ to t). Flow utility from consumption and bequests are $u(C) = \frac{C^{1-\sigma}-1}{1-\sigma}$ and $v(B) = \theta \frac{(\kappa+B)^{1-\sigma}}{1-\sigma}$, respectively.

The household faces exogenous risk in its income and housing costs. Before retirement, income, rents, and home prices evolve stochastically according to

$$x_t = \alpha + \gamma_t + \Gamma x_{t-1} + \eta_t, \tag{C.2}$$

[21]This value is high but is consistent with the 'naive' ratio obtained by dividing median home value by median rent. Of course, this does not take into account that the median owner-occupied housing unit is larger than the median rental unit; similarly, the median owner has much higher income than the median renter. Recent research using rental and transaction prices for the same property starting in 2000 finds that the average price-rent ratio is about 10; see Begley et al. (2019). As mentioned above, we instead use the implied expected carrying costs for homeowners to make rental and owner-occupied housing units comparable.

[22]The choice of decade-level timing is motivated by the available housing and earnings data.

[23] $\prod_{s=t+1}^{t}(1-\delta_s) = 1$; i.e., the empty product is one.

where $x_t \equiv \log(Y_t, R_t, P_t)'$. The coefficients in α and Γ are estimated using the VAR described in Section B.3. The γ_t term allows for realistic life-cycle profiles of earnings, with a peak in the 40s. The stochastic shocks η_t are normally distributed, $\eta_t \sim iid\ \mathcal{N}(0, \Sigma)$, where the covariance matrix Σ is estimated based on the residuals of the VAR. Our main VAR specification includes CZ and time fixed effects (α and γ_t). It is therefore identified by within-location changes over time and so excludes national shocks. This leads us to understate the risk facing households, but it reduces concerns about predictable long-run changes that might not reflect risk such as a trend in technological progress or changes in the quality or "efficiency units" of housing. During retirement (age decade 60 and onwards), income is constant at half of end-of-career earnings and no longer interacts with rents and home prices (covariances in Γ and Σ are set to zero).

The budget constraint is

$$\begin{aligned} A_{t+1} &= (1+r)\left[A_t + Y_t - \text{Rent}_t R_t - \text{Buy}_t P_t\,(1+\zeta_p) + \text{Sell}_t P_t\,(1-\zeta_s) + \theta_t - C_t\right] \\ &\geq -\text{Own}_t \phi_t P_t (1+r). \end{aligned} \tag{C.3}$$

Non-housing assets, A_t earn risk-free return r. Rent_t and Own_t are indicators of whether the household is currently renting or owning ($\text{Rent}_t + \text{Own}_t = 1\ \forall t$). Buy_t and Sell_t are indicators of whether the household (exogenously) bought or sold a home in period t. Purchasing a home costs P_t plus a proportional transaction fee ζ_p. Selling a home nets P_t less a proportional transaction fee ζ_s.[24] Consumption spending, c_t, is constrained by a loan-to-value constraint: The household can accumulate debt only up to fraction ϕ_t of the value of its home: $\text{Own}_t \phi_t P_t$. This embeds the standard assumption that there is no uncollaterized borrowing, so renters cannot borrow: $A_{t+1} \geq 0$ if $\text{Own}_t = 0$. Transfers from means-tested programs, θ_t are such that the household can consume $\bar{c}$ while meeting its borrowing constraint with equality, $A_{t+1} = -\text{Own}_t \phi_t P_t (1+r)$. Any non-negative assets remaining at death are bequeathed, including both liquid wealth and any net proceeds from the sale of a home, $B_{t+1} = \max\left\{0, A_{t+1} + \text{Own}_t\,(1-\zeta_s)\,P_{t+1}\right\}$. Table A.1 shows values of all of the parameters of the model, which follow those in Davis and Van Nieuwerburgh (2015) wherever possible. Appendix C.4 contains details about our solution procedure.

C.2 Life Cycle Model Welfare Calculations

Table A.2 shows the means and standard deviations of the lifetime present value of simulated income, housing costs, and consumption of an owner versus a renter for each of four specifications: no risk, earnings risk only, earnings and housing cost risk as estimated in the data, and risk as estimated except with the correlations between earnings and housing costs set to zero.[25] The statistics are based on 10,000 simulated lifetimes of both a renter and an owner. Simulated individuals begin their life cycles in their twenties, with no assets and with (Y_t, R_t, P_t) equal to its unconditional mean.

The results are quite similar to those based on the observed empirical distribution in Section 5. Risk in wages, rents, and home prices leads to substantial risk in consumption, and much more so for owners than for renters. For a household without any exposure to housing cost risk (e.g., a household that fully insured with actuarially fair housing cost insurance), the standard deviation of lifetime consumption is $105,000. A homeowner—even the buy-and-hold-for-life owner considered in these calculations—is exposed to somewhat more risk. The standard deviation of an owner's lifetime consumption is $113,000, which is about 9% of their lifetime net income of $1.3m (defined

[24] Using notation from Section 2, $P_t^{buy} = P_t\,(1+\zeta_p)$ and $P_t^{sell} = P_t\,(1-\zeta_s)$.

[25] To isolate differences in risk across specifications, we adjust the γ_t terms in order to hold fixed the expected present value of net income across specifications. Details are in Appendix C.3.

as the expected present value of income net of housing cost) and is about 8% greater than someone without exposure to housing cost risk. A renter is exposed to significantly less consumption risk. The standard deviation of a renter's lifetime consumption is $84,000, which is about 7% of their lifetime net income of $1.2m and is about 20% less than someone without exposure to housing cost risk.

Next, we turn to measuring the welfare cost of the different risk exposures of owners versus renters. Let $V_0(A_0; \text{risk})$ and $V_0(A_0; \text{no risk})$ denote initial-period value functions in specifications with and without risk, respectively, with initial assets A_0. A compensating variation (willingness to pay) measure of the welfare cost of a given risk exposure, CV, is defined implicitly by

$$V_0(A_0 - CV; \text{no risk}) = V_0(A_0; \text{risk}). \tag{C.4}$$

Column 1 of Table A.3 shows the results. Whereas a household without exposure to housing cost risk is willing to pay 2.7% of its lifetime net income to eliminate risk, a renter is willing to pay 2.1% and an owner is willing to pay 2.9%. Even for the buy-and-hold-for-life owner considered here, exposure to housing cost risk is costly. This cost, about 0.5% of lifetime net income, comes largely from the positive correlation between house prices and earnings (i.e., the portfolio effect) rather than the direct effect of house price risk (though it is important to note that the portfolio effect likely increases in importance relative to the direct effect as the household's horizon increases). By contrast, the renter's housing cost risk exposure is valuable, worth about 0.6% of lifetime net income, or about $6,000 given lifetime net income of $1m. The overall risk exposure advantage of renting over owning is therefore about 1.1% of lifetime net income, or about $11,000 for a household with $1m of lifetime net income.

The key role of the correlations between earnings and housing costs can be seen by comparing the main results to those from a model in which these correlations are set to zero (column 2). With uncorrelated earnings and housing costs, rent risk is quite costly: A renter is willing to pay 0.7% of lifetime net income to eliminate it. An owner's housing risk, on the other hand, is less costly: A buy-and-hold-for-life owner is willing to pay just 0.2% of lifetime net income to eliminate its housing cost risk. So if earnings and housing costs were uncorrelated, the advantage of owning over renting would be roughly 0.5% of lifetime net income, or about $5,000. This quantifies the mechanism emphasized in Sinai and Souleles (2005): With uncorrelated earnings and housing costs, owning is safer than renting for households with long enough horizons in a home, since the value of an owner's hedge against rent risk exceeds the cost of its exposure to home price risk. But with the strong positive correlations between earnings and housing costs observed empirically, this conclusion reverses: Renting is safer than owning, even for the buy-and-hold-for-life owner.

C.3 Log-Normal Correction

We hold fixed the expected present value of net income across specifications by including an age-varying log-normal correction term in the model. Specifically, the γ_t term in (C.2) is constructed from two additive components: $\gamma_t = \gamma_t^{LN} + \gamma_t^{LC}$, where γ_t^{LC} is the mean earnings at every age decade (based on data), and γ_t^{LN}—the focus of this section—is

$$\gamma_t^{LN} = -0.5 diag(\text{Var}_0(x_t)). \tag{C.5}$$

Here, $diag(\cdot)$ is the operator that extracts the diagonal of a matrix into a vector, and $\text{Var}_0(x_t)$ is the variance of future x_t conditional on information at time 0, computed iteratively as

$$\begin{aligned}
&\text{Var}_0(x_0) = 0 \quad \text{(no initial uncertainty)},\\
&\text{Var}_0(x_t) = \Gamma \text{Var}_0(x_{t-1})\Gamma' + \Sigma.
\end{aligned}$$

C.4 Details of solution to the life cycle model

The recursive version of the problem is

$$V_t(A_t, x_t) = \max_{C_t, A_{t+1}} \left\{ u(C_t) + \tilde{V}_t\left(A_{t+1}, x_t\right) \right\} \tag{C.6}$$

subject to the budget and borrowing constraints (C.3), the consumption floor, and the bequest equation $B_{t+1} = \max\left\{0, A_{t+1} + \text{Own}_t\left(1-\zeta_s\right)P_{t+1}\right\}$. The (discounted) expected continuation value function $\tilde{V}_t\left(A_{t+1}, x_t\right)$ in the last possible period of life T is

$$\tilde{V}_T\left(a_{T+1}, x_T\right) = \beta E_T\left[\delta_{T+1} v(\max\left\{0, a_{T+1} + \text{Own}_t P_{T+1}\right\})\right], \tag{C.7}$$

and for all other periods is

$$\begin{aligned}
\tilde{V}_t\left(A_{t+1}, x_t\right) &= \beta E_t\left[(1-\delta_{t+1})V_{t+1}\left(A_{t+1}, x_{t+1}\right)\right]\\
&\quad + \beta E_t\left[\delta_{t+1} v(\max\left\{0, A_{t+1} + \text{Own}_t P_{t+1}\right\})\right]\\
&\approx \beta E_t\left[(1-\delta_{t+1})V_{t+1}\left(A_{t+1}, x_{t+1}\right)\right]\\
&\quad + \delta_{t+1}\tilde{V}_T\left(A_{t+1}, x_t\right),
\end{aligned} \tag{C.8}$$

where in the approximation, $\tilde{V}_T\left(A_{t+1}, x_t\right) = \tilde{V}_T\left(a_{T+1}, x_T\right)|_{a_{T+1}=A_{t+1}, x_T=x_t}$ (i.e. it is to be viewed as a function evaluated at a point). The approximation is exact if Γ and Σ did not depend on age, and is useful for computations because we only need to compute the integral once and store it.

We solve the model recursively backwards using the endogenous grid method, implemented based on the procedure described in Fella (2014) and similar to that in Ameriks, Briggs, Caplin, Shapiro and Tonetti (2015). Starting from the last time period, we do the following:

1. For every A_{t+1}^{target}, x_t^i, and tenure status (rent or own) on a grid, we compute the derivative of the expected continuation value function with respect to the asset choice.[26]

   The functional form of the derivative follows from straightforward differentiation of (C.7) and (C.8):

$$\frac{\partial \tilde{V}_T\left(a_{T+1}, x_T\right)}{\partial a_{T+1}} = \beta v'\left(a_{T+1}\right) \quad \text{for renters,} \tag{C.9}$$

$$\frac{\partial \tilde{V}_T\left(a_{T+1}, x_T\right)}{\partial a_{T+1}} = \beta E_T\left[v'\left(a_{T+1} + P_{T+1}\right)\mathbb{1}\left[P_{T+1} \geq -a_{T+1}\right]\right] \quad \text{for owners,} \tag{C.10}$$

[26] We use superscripts to indicate specific values, usually values on grid points. Superscript i are for objects that do not cause confusion, superscript *target* indicates target savings amounts, and superscript *implied* indicates initial asset amounts that are implied by other values on grids.

$$\frac{\partial \tilde{V}_t\left(A_{t+1}, x_t\right)}{\partial A_{t+1}} = \beta E_t\left[\left(1-\delta_{t+1}\right)\frac{\partial}{\partial A_{t+1}} V_{t+1}\left(A_{t+1}, x_{t+1}\right)\right] + \delta_{t+1}\frac{\partial \tilde{V}_T\left(A_{t+1}, x_t\right)}{\partial a_{T+1}}. \tag{C.11}$$

(C.9) and (C.10) are easy to compute on the grid. Computation of $\frac{\partial \tilde{V}_t\left(A_{t+1}^{target}, x_t^i\right)}{\partial A_{t+1}}$ using (C.11) requires the derivative of the value function from Step 7. Furthermore, because expectations over continuous random variables are involved, evaluations of $\frac{\partial}{\partial A_{t+1}} V_{t+1}\left(A_{t+1}^{target}, x_{t+1}\right)$ are performed off the x_t^i grid from Step 7.[27] This is done using a 3-dimensional linear interpolant.[28]

2. For every A_{t+1}^{target}, x_t^i and tenure status on a grid, we compute the assets $A_{t+1}^{implied}$ implied by the first order conditions for interior solutions (i.e. ignoring the borrowing constraint and the consumption floor):

$$A_t = \left(u'\right)^{-1}\left(\left(1+r\right)\frac{\partial \tilde{V}_t\left(A_{t+1}, x_t\right)}{\partial A_{t+1}}\right) - Y_t + \frac{1}{1+r}A_{t+1} + \text{Rent}_t R_t,$$

3. If $t < T$, we do the following steps. For every x_t^i and tenure status on a grid:
   (a) We split the grid for A_{t+1}^{target} into two regions to use the method in Fella (2014): a definitely concave and a possibly nonconcave region (henceforth concave and nonconcave regions). The split is done by locating the first and last intervals for which $\tilde{V}_t\left(A_{t+1}^{target}, x_t^i\right)$ is strictly decreasing in A_{t+1}^{target}. These two intervals constitute the concave region; the remaining points are the nonconcave region.
   (b) The concave region consists of A_{t+1}^{target}'s on the grid in which the maximization problem ignoring the no-borrowing constraint and government transfers is definitely concave. For these A_{t+1}^{target}'s, the $A_t^{implied}$ from Step 2 are the unique solution and are saved in a solution array.
   (c) In the nonconcave region, we check if the $A_t^{implied}$ from Step 2 actually maximize utility for the associated candidate A_{t+1}^{target} before saving in the solution array. To do this, we fix $A_t^{implied}$, and compute the utility levels implied by each of the A_{t+1}^{target} in the nonconcave region. If the candidate A_{t+1}^{target} does not give the highest utility among all the A_{t+1}^{target}'s, we discard this point.
4. For $t = T$, we identify the nonconcave region in a simpler way. Taking two derivatives with respect to a_{T+1} of the last period's maximization problem (ignoring the borrowing constraint and consumption floor) shows that the problem is concave if and only if $\frac{\partial a_T}{\partial a_{T+1}} \geq 0$. Hence, the solution array for the last period comprises $A_t^{implied}$ that are increasing in A_{t+1}^{target}.
5. For every x_t^i and tenure status on a grid, we invert the A_{t+1} to A_t mapping by interpolating $\left(A_t^{implied}, A_{t+1}^{target}\right)$ to obtain, on a grid of A_t's, the choice of A_{t+1}.[29] This gives the first order condition-derived interior solution, which at this step is also the candidate optimal solution $A_{t+1}^{candidate}$.

[27] Note that we choose the A_{t+1}^{target} grid to coincide with the A_t^i grid from Step 7. This removes the need to interpolate in one more dimension.

[28] Linear interpolation is used for speed, and because its result is invariant to the choice of which dimension to interpolate first.

[29] Piecewise cubic hermite interpolating polynomial (PCHIP) interpolation is used.

6. We check if the interior solution is the correct one, i.e. update the solution if the individual is borrowing constrained or should be using the consumption floor.

   (a) For every A_t^i, x_t^i and tenure status on a grid, we check if the borrowing constraint is violated. If violated, we set $A_{t+1}^{candidate}$ to the constraint value A_{t+1}^{BC}.

   (b) For every A_t^i, x_t^i and tenure status on a grid, we replace $A_{t+1}^{candidate}$ with A_{t+1}^{BC} if the latter is preferred. This step is done ignoring the consumption floor.

   (c) For every A_t^i, x_t^i and tenure status on a grid, we check if the individual prefers to use the consumption floor. This sets $A_{t+1}^{candidate}$ to A_{t+1}^{BC} if so.

7. The above steps give the policy functions (liquid savings, consumption, and usage of consumption floor) and the value functions, evaluated on a grid of A_t^i,x_t^i, and tenure status. In preparation for the next step of the loop, we compute the derivative of the value function $\frac{\partial}{\partial A_{t+1}} V_t\left(A_t^i, x_t^i\right)$. This derivative is computed analytically; application of the envelope theorem on (C.6) gives

$$\frac{\partial}{\partial A_t} V_t\left(A_t, x_t\right) = u'\left(C_t^*\right)$$

   where C_t^* is the consumption policy function at time t.

8. We then repeat Steps 1 to 7 for the previous time period.

Table A.1 – Life cycle model parameter values

Symbol	Parameter	Value
	Utility parameters	
σ	CRRA utility parameter	3
κ	Bequest parameter[a]	\$410,000
θ	Bequest parameter[a]	13.8
	Stochastic processes and life-cycle objects	
t	Periods of life	20s to 90s (8 decades)
$\mathbf{\Sigma}$	Covariance matrix of stochastic shocks before age 60	$\begin{bmatrix} 0.0091 & 0.0113 & 0.0147 \\ 0.0113 & 0.0339 & 0.0346 \\ 0.0147 & 0.0346 & 0.0779 \end{bmatrix}$
$\mathbf{\Sigma}^r$	Covariance matrix of stochastic shocks from age 60	$\begin{bmatrix} 0 & 0 & 0 \\ 0 & 0.0339 & 0.0346 \\ 0 & 0.0346 & 0.0779 \end{bmatrix}$
ρ	VAR coefficient before age 60	$\begin{bmatrix} 0.95 & 0.04 & -0.02 \\ 0.07 & 0.92 & 0.01 \\ -0.15 & 0.13 & 0.88 \end{bmatrix}$
ρ^r	VAR coefficient from age 60	$\begin{bmatrix} 1 & 0 & 0 \\ 0 & 0.92 & 0.01 \\ 0 & 0.13 & 0.88 \end{bmatrix}$
α	VAR constant	$\begin{bmatrix} 0.35 & -0.12 & 1.87 \end{bmatrix}'$
δ_t	Probability of dying before each age[b]	Age 30: 1.1%; 40: 1.5%; 50: 2.6%; 60: 6.0%; 70: 12.1%; 80: 26.6%; 90: 59.0%; 100: 100%
	Ratio of earnings in each decade of life to average earnings	20s: 0.83; 30s: 1.02; 40s: 1.07; 50s: 1.01
	Ratio of retirement earnings to 50s earnings	0.5
	Other parameters	
$1+r$	Decadal rate of return on savings	1.34
β	Decadal discount factor	0.74
$\underline{c}$	Consumption floor	\$50,000 per decade
ζ_p, ζ_s	Transaction costs on purchasing or selling a house	3% of house value
ϕ_t	Borrowing constraint parameter for homeowners	80% at $\delta_{t+1} = 0$, decreases linearly to 0% at $\delta_{t+1} = 1$

[a] Estimates from Lockwood (2018) rescaled to reflect decadal time periods.
[b] Based on the Social Security Administration's Actuarial Life Table for 2016.

Table A.2 – Life-cycle model statistics

	Risk in model							
	Riskless	Earnings risk only	All risks, uncorre-lated	All risks, corre-lated	Riskless	Earnings risk only	All risks, uncorre-lated	All risks, corre-lated
	Values for renters				*Values for owners*			
PV of income	1,513	1,514	1,511	1,514	1,513	1,514	1,511	1,514
	[0]	[121]	[123]	[125]	[0]	[121]	[123]	[125]
PV of housing costs	328	328	329	327	227	227	227	227
	[0]	[0]	[50]	[52]	[0]	[0]	[21]	[24]
PV of consumption	1,114	1,112	1,106	1,111	1,207	1,205	1,200	1,202
	[0]	[105]	[111]	[84]	[0]	[104]	[106]	[113]
PV of bequests	71	77	78	77	80	85	87	88
	[0]	[19]	[20]	[17]	[0]	[20]	[24]	[29]

Notes: Table shows the means and standard deviations of the present value (PV) of various model outputs, based on 10,000 simulated individuals. Each cell shows the mean, followed by the standard deviation in brackets. All values are in thousands of 1990 dollars. PVs incoporate the probability of dying in each time period. Housing costs for renters are rents they pay each period. Housing costs for owners are the initial house price less the discounted house price at the time of death, with a 3% transaction cost imposed in each of the two periods.

Table A.3 – Life-cycle model welfare results

	Estimated risks (1)	Correlations set to 0 (2)
Panel A: Renter's WTP to eliminate specified risks (% of EPVNY)		
Exposure to income risk only	2.7%	2.7%
Exposure to all risks	2.1%	3.4%
Value of renter's hedge against income risk	**0.6%**	
Panel B: Owner's WTP to eliminate specified risks (% of EPVNY)		
Exposure to income risk only	2.3%	2.3%
Exposure to all risks	2.9%	2.5%
Cost of owner's exposure to home price risk	**0.5%**	
Panel C: Total risk advantage of renting over owning (% of EPVNY)		
Total value	**1.1%**	

Notes: Ex ante willingness to pay (WTP) to eliminate specified risks by a household beginning its life in a location in which the income, rent, and home price process is at its unconditional mean. All are expressed as a percentage of the expected present value of income net of housing costs, $EPVNY$. *Exposure to income risk only* refers to models in which only income is risky. *Exposure to all risks* refers to models in which income, rents, and home prices are risky. In column 1, these risks are as estimated empirically. In column 2, the correlations between housing prices and income are set to zero. *Value of renter's hedge against income risk* is the renter's WTP to eliminate income risk less its WTP to eliminate all risks, i.e., the extent to which exposure to rent risk reduces the welfare cost of its overall risk exposure. *Cost of owner's exposure to home price risk* is the owner's WTP to eliminate all risks less its WTP to eliminate income risk, i.e., the extent to which exposure to home price risk increases the welfare cost of its overall risk exposure. *Total risk advantage of renting over owning* is the sum of the value of the renter's hedge against income risk and the cost of the owner's exposure to home price risk.

Table A.4 – Correlation matrix of lifetime PVs

		Y	R	$-P_T^{net}$		
				a.	b.	c.
Y		1.00				
R		0.74	1.00			
$-P_T^{net}$	a.	-0.22	-0.49	1.00		
	b.	0.45	0.56	-0.65	1.00	
	c.	-0.22	-0.49	1.00	-0.65	1.00

Figure A.1 – County-Level Correlation of Property Tax Rate Changes and Income Growth

(a) property tax rate/income: all variation

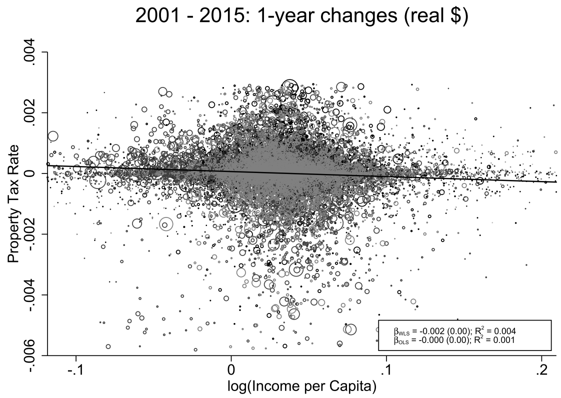

(b) property tax rate/income: location variation

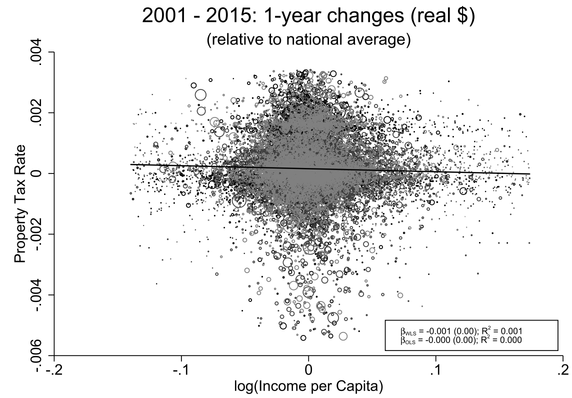

(c) property tax rate/income: all variation

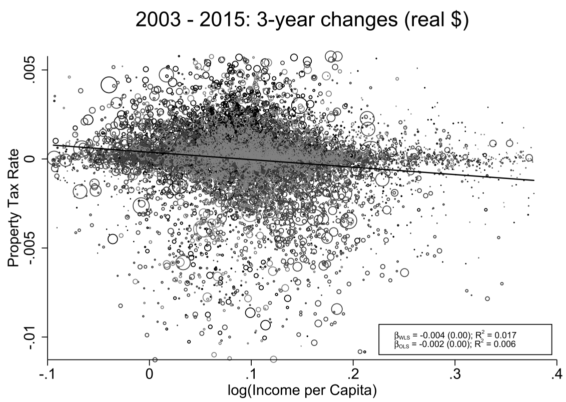

(d) property tax rate/income: location variation

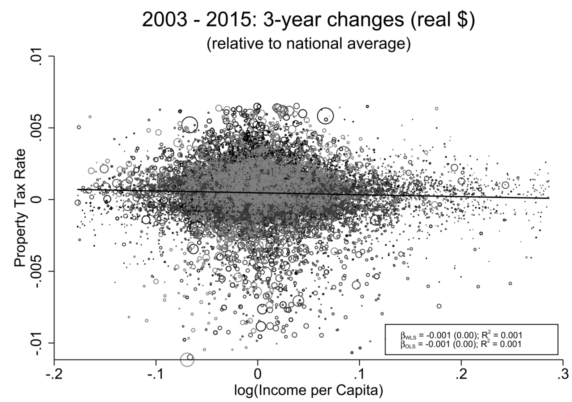

(e) property tax rate/income: all variation

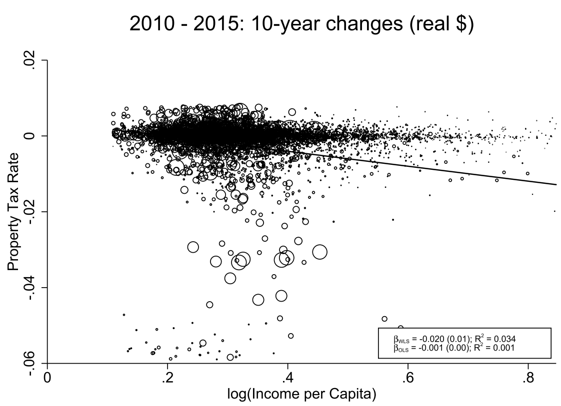

(f) property tax rate/income: location variation

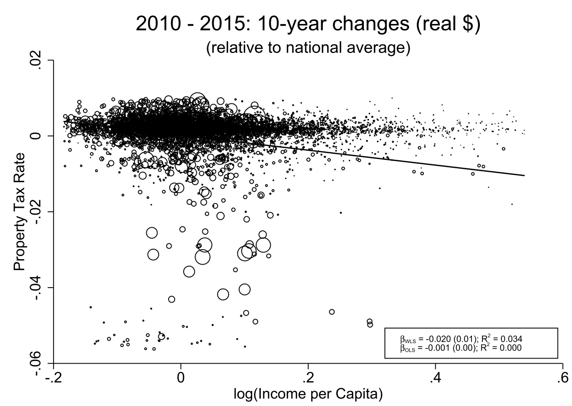

Figure A.2 – National-Level Correlation of Full-Time Wage Income and Housing Costs Growth

(a) rent/income: average hourly earnings

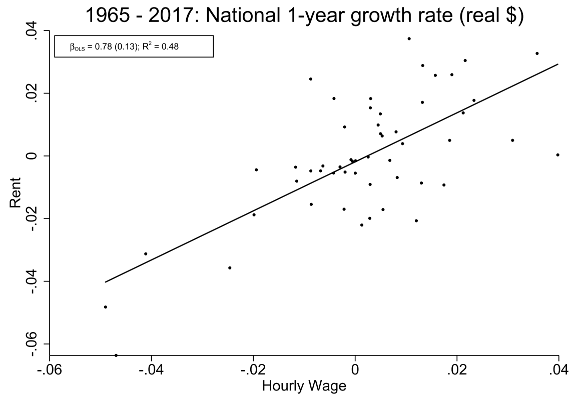

(b) price/income: average hourly earnings

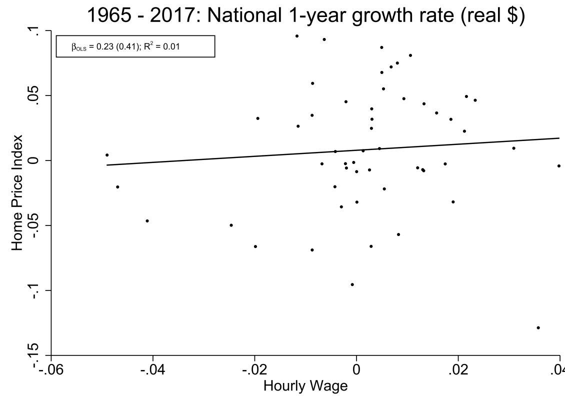

(c) rent/income: usual weekly earnings

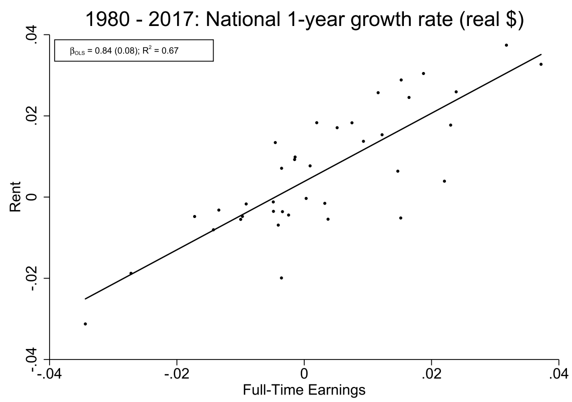

(d) price/income: usual weekly earnings

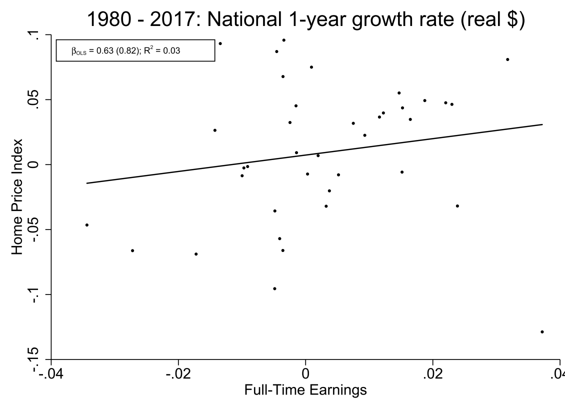

(e) rent/income: Decennial Census data

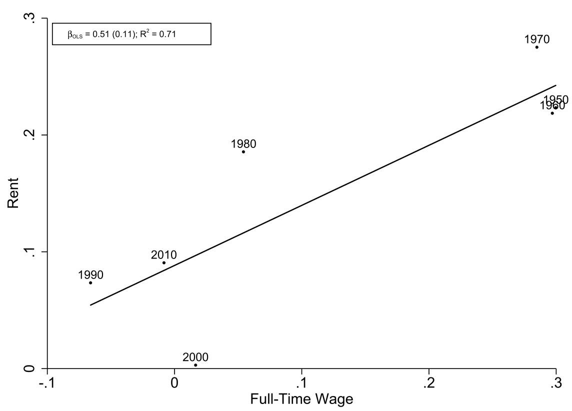

(f) price/income: Decennial Census data

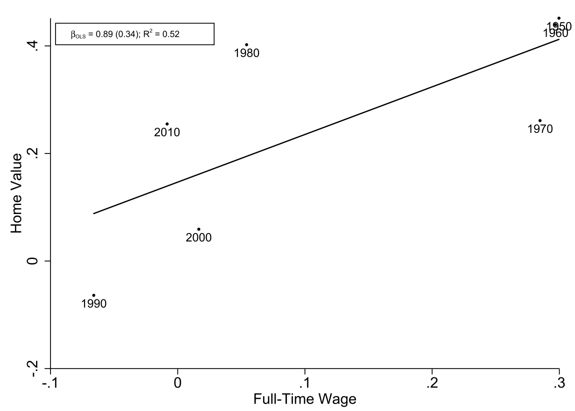

Figure A.3 – National-Level Index of Hourly Wage, Full-Time Salary, Rent and Home Price

(a) BLS data

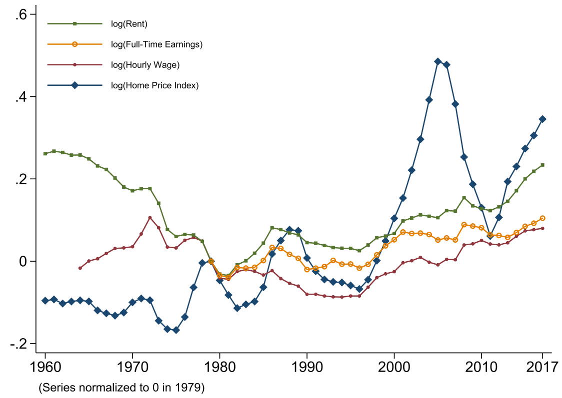

(b) VAR(2)-based series vs. data: income

(c) VAR(2)-based annual series vs. data: rent

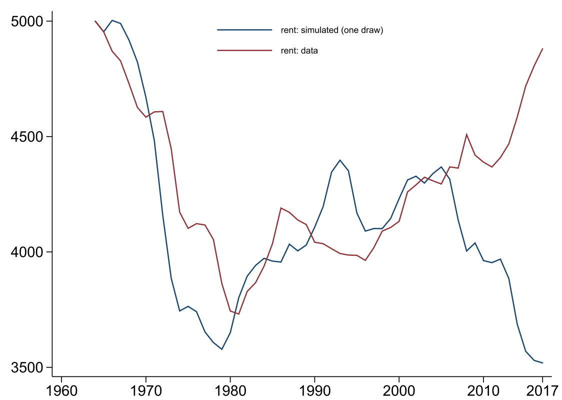

(d) VAR(2)-based annual series vs. data: price

(e) rent/income: simulated data

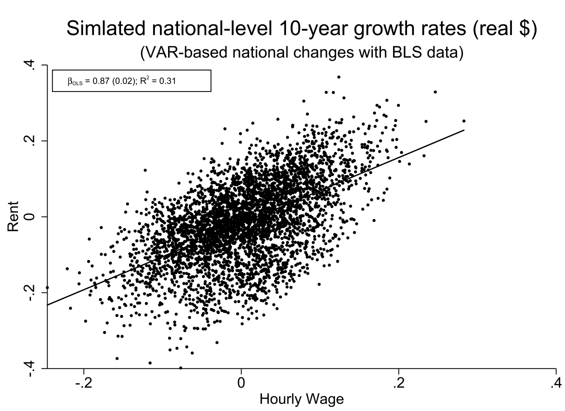

(f) price/income: simulated data

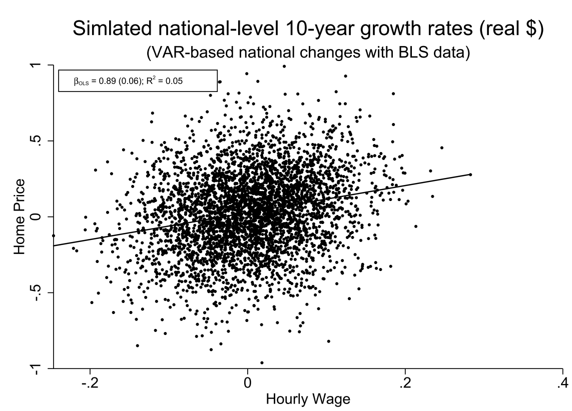

Figure A.4 – (Non-)Correlation of Income, Social Security, and Utility Costs

(a) Social Security/income

(b) electricity cost/income

(c) fuel cost/income

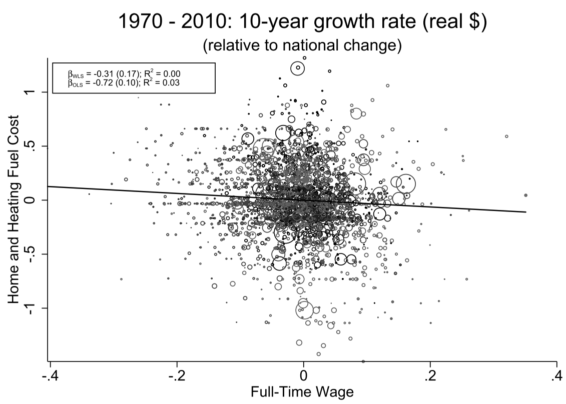

(d) gas cost/income

(e) water cost/income

(f) all utilities/income

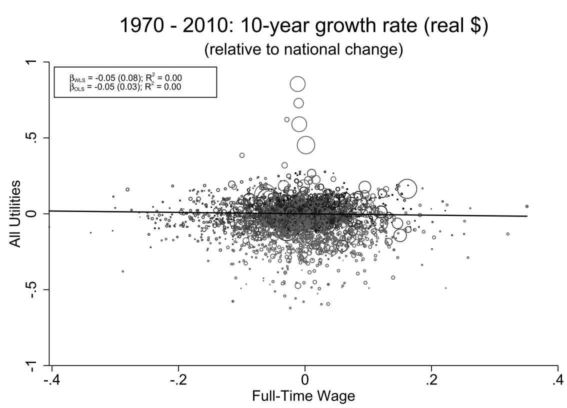

Notes. Social Security is a federal program without local cost of living adjustments, hence it should not be strongly related to local wage changes. Similarly, electricity is supplied competitively on a national grid and prices are mostly determined at the national level.

Figure A.5 – Robustness to using Gross Rent instead of Contract Rent

(a) gross rent/income: all 10-year variation

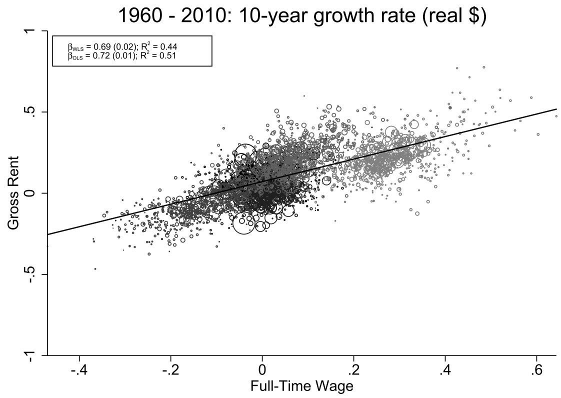

(b) rent/income: all variation

(c) gross rent/income: location variation

(d) rent/income: location variation

(e) gross rent/income: long-run variation

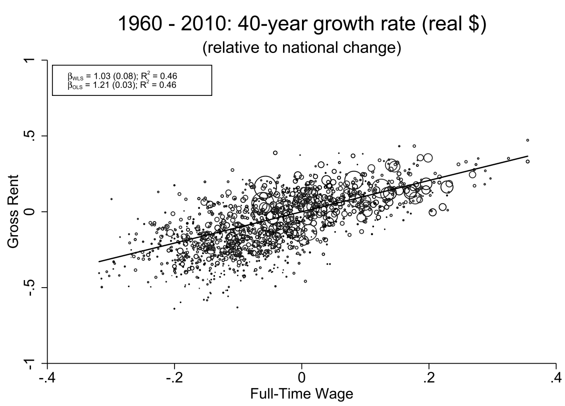

(f) rent/income: long-run variation

Figure A.6 – Robustness to Dropping 2000s

(a) rent/income: all 10-year variation

(b) price/income: all 10-year variation

(c) rent/income: location variation

(d) price/income: location variation

(e) rent/income: long-run variation

(f) price/income: long-run variation

Figure A.7 – Robustness to Dropping 1940s

(a) rent/income: all 10-year variation

(b) price/income: all 10-year variation

(c) rent/income: location variation

(d) price/income: location variation

(e) rent/income: long-run variation

(f) price/income: long-run variation

Figure A.8 – Robustness to Dropping 1940s and 1950s

(a) rent/income: all 10-year variation

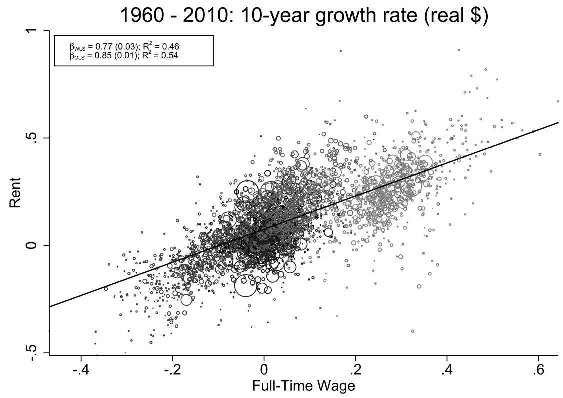

(b) price/income: all 10-year variation

(c) rent/income: location variation

(d) price/income: location variation

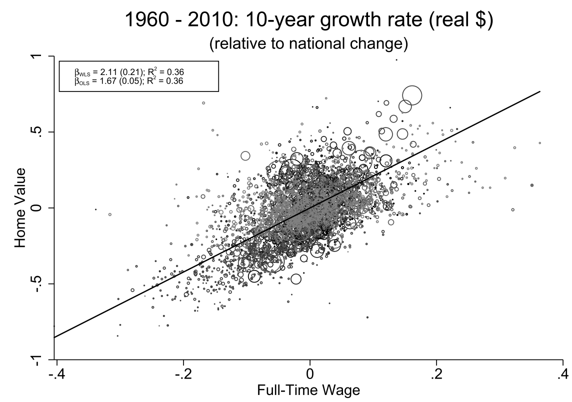

(e) rent/income: long-run variation

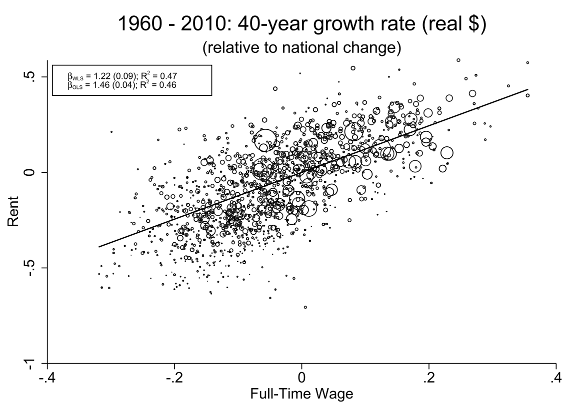

(f) price/income: long-run variation

Figure A.9 – Robustness to using Counties

(a) rent/income: all 10-year variation

(b) price/income: all 10-year variation

(c) rent/income: location variation

(d) price/income: location variation

(e) rent/income: long-run variation

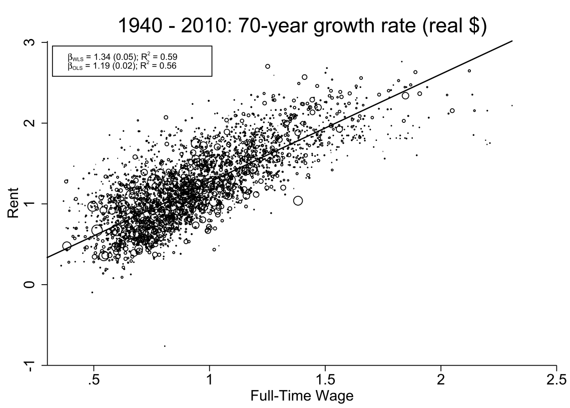

(f) price/income: long-run variation

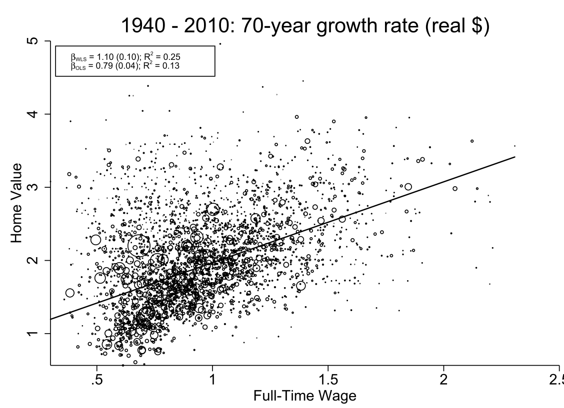

Figure A.10 – Robustness to Mobility (2) – Controlling for Population Changes

(a) rent/income: all 10-year variation

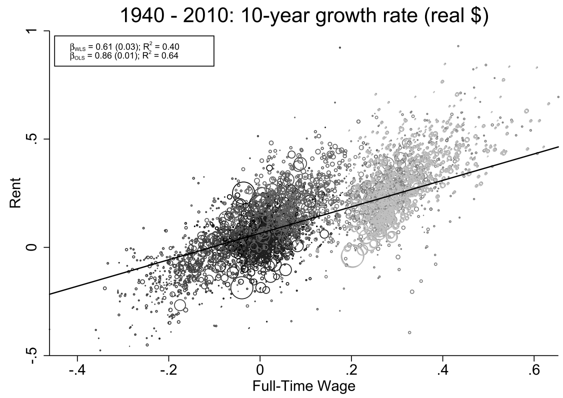

(b) price/income: all 10-year variation

(c) rent/income: location variation

(d) price/income: location variation

(e) rent/income: long-run variation

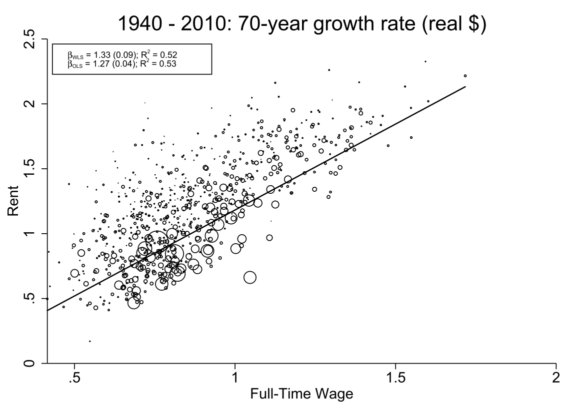

(f) price/income: long-run variation

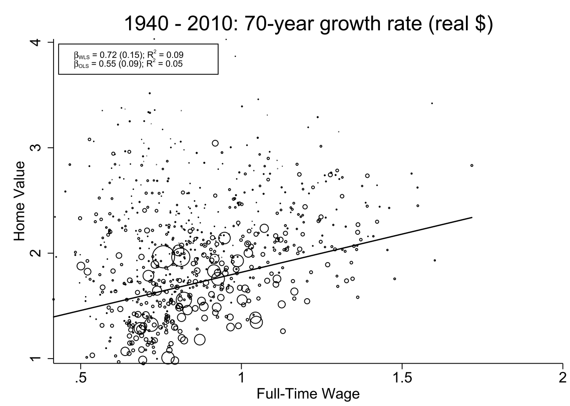

Figure A.11 – Robustness to Mobility (1) – Dropping Movers

(a) rent/income: all 10-year variation

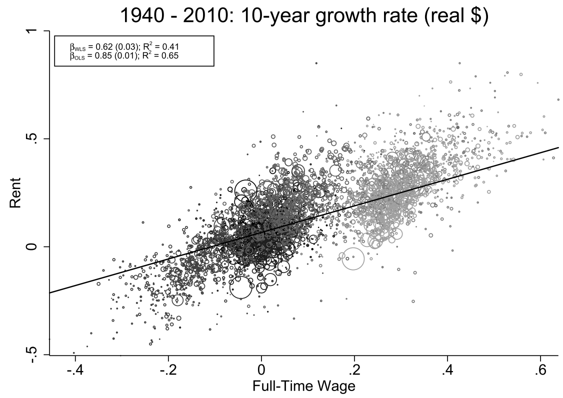

(b) price/income: all 10-year variation

(c) rent/income: location variation

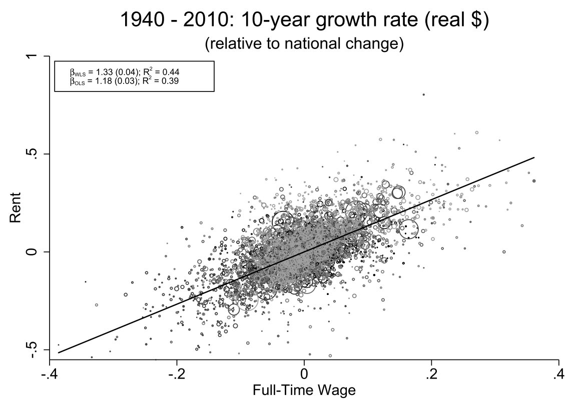

(d) price/income: location variation

(e) rent/income: long-run variation

(f) price/income: long-run variation