The Riskiness of Owning Versus Renting Housing*

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

Preliminary and incomplete: Please do not circulate without permission.
Click here for the most recent version: https://bit.ly/36IPr0s

April 3, 2022

Abstract

Homeowners and renters have mirror-image exposures to the substantial risks in housing prices. The welfare costs of these exposures depend 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, earnings risk is hedged by renting and exacerbated by owning. These interactions with earnings risk are sufficiently strong that, contrary to widely-held views, for many households owning is significantly riskier than renting, regardless of the time horizon in the home. This highlights an important cost of owner-occupied housing and the many major policies that encourage it.

JEL Classification: D12, D14, E21, G11, G51
Keywords: housing tenure, income risk, rent risk, home price risk, local labor markets

*We thank Jeff Campbell, Eric Chyn, Kerem Cosar, Leland Farmer, Greg Howard, Matt Notowidigdo, Ed Olsen, Isaac Sorkin, and seminar participants at the Chicago Fed, Claremont McKenna, the Empirical Macro Workshop, Bern, Geneva, IFS (UCL/LSE), Luzern, Munich, Northwestern, NYU, Salzburg, UVA, and William & Mary for helpful discussions. We gratefully acknowledge financial support from the Bankard Fund for Political Economy (Lockwood) and Swiss National Science Foundation grant No. 100018_197802 (Kueng).

†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, Swiss Finance Institute, and Centre for Economic Policy Research. E-mail: lorenz.kueng@usi.ch.

¶National University of Singapore, Department of Strategy and Policy, 15 Kent Ridge Drive, Singapore 119245. E-mail: ong.pinchuan@nus.edu.sg.

1 Introduction

Housing is a major consumption good and an important asset, comprising about one third of both household expenditures and household net assets. A household's exposure to the substantial risks in "housing prices" (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 during their stay. So while owning exposes a household to significant asset price risk, it hedges the significant risk in rents it would otherwise face as a renter (Sinai and Souleles, 2005). Whether due to this or other considerations, the prevailing view is that for many households owning is safer than renting.[1]

In this paper, we find instead that for many households, owning is significantly *riskier* than renting regardless of their horizon in their home. The key proximate driver is the strong positive correlations between market-level housing prices and wages: Higher wages are strongly associated with higher home prices and higher rents. This creates strong, opposite-signed correlations between household earnings—a major source of hard-to-diversify risk—and the housing costs of owners versus renters. An owner's housing costs are negatively correlated with wages (since wages are positively correlated with the price at which the owner can sell its home), so an owner's housing exposure exacerbates wage risk. A renter's housing costs are strongly positively correlated with wages (since wages are positively correlated with the rents a renter must pay to rent its home), so a renter's housing exposure provides a valuable hedge against wage risk. Although for households with long enough horizons in a home, owning tends to lead to less risk in their *housing costs*, renting tends to lead to less risk in their portfolios as a whole—and so in their consumption and welfare.

The key underlying mechanism is the substantial location-specific risk in wages and housing prices, which gives rise to strong positive relationships between location-specific changes in wages and housing prices. To fix ideas, consider the experiences of residents of Detroit and Seattle late in the 20th century. Detroit's decline imposed much greater losses on local homeowners than local renters, as owners experienced substantial declines in the value of their homes on top of their depressed earnings, whereas renters received partial compensation for their depressed earnings in the form of lower rents. Similarly, Seattle's boom conferred much greater benefits on local homeowners than lo-

[1] This view is widespread in the economics literature, 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 et al. 2018). For example, in a nationally representative survey of more than 50,000 households, Adelino et al. (2018) find that 71% of U.S. households believe that housing is a "safe" investment, versus 18% for stocks and 55% for government and corporate bonds.

cal renters, as owners experienced substantial gains in the value of their homes on top of their increased earnings, whereas renters experienced increases in their housing costs that partially offset their increased earnings. Ex ante, for households who do not know if their city will go the way of Detroit or Seattle, renting provides a valuable hedge against the substantial location-specific risk in wages, whereas owning "doubles down" on it.

We begin by using over 70 years of data on local markets in the U.S. (commuting zones) to characterize the implications of the volatility of wages, rents, and home prices for the volatility of owners' versus renters' housing costs and net income (income net of housing costs). This analysis produces three main findings. First, the substantial volatility of housing costs, which arises from the substantial volatility of rents and home prices, is minimized by renting if a household has a shorter horizon in a home but by owning if a household has a longer horizon in a home (longer than 15–20 years or so). So were it not for other risk exposures, this finding would be consistent with the prevailing view that owning is safer than renting for households with long enough horizons in a home. But even in this case, the horizon at which owning becomes safer is longer than typical lengths of time between moves.

Second, changes in home prices and rents are strongly positively correlated with changes in wages: Locations that experience faster growth in home prices and rents during a particular time period are much more likely to have experienced faster growth in wages as well, over all horizons from 10 to 70 years. This suggests that for working households there are strong interactions between risk in housing costs and risk in earnings, where renting hedges earnings risk—with lower rents "compensating" for lower wages and higher rents partially offsetting higher wages—and owning exacerbates it.

Third, the interaction with earnings risk is sufficiently strong that for working households, renting *reduces* overall risk in net income—net income is less variable by renting than it would be even with perfect housing cost insurance—whereas owning increases risk in net income. These hold regardless of the time horizon in a home. Hence, for working households, renting is significantly less risky than owning regardless of the time horizon in the home.

We then investigate, using a variety of approaches, the implications of these strong empirical regularities for the welfare consequences of the risk aspect of owning versus renting housing. Each approach points to the same conclusion: For most households, renting confers a sizable risk benefit whereas owning involves a significant risk cost. For example, we estimate that even for a household that never moved, the reduction in risk from switching from always owning to always renting is likely to be worth on the order of $10,000 in lifetime present value.

The sizable risk benefit for many households of renting over owning reflects two reinforcing factors. First, renters avoid the substantial welfare cost of owners' sales price risk. This cost is substantial because home prices are not only highly variable—much more so than rents—but positively correlated with wages.[2] Second and more important, renting provides a valuable hedge against the substantial income risk from risk in market-level wages. Our estimates suggest that 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, so exposure to rent risk *reduces* overall net income and consumption risk. 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 costs.

The crucial role of a household's exposure to wage risk highlights systematic heterogeneity in the risk consequences of owning versus renting housing. Households with less exposure to wage risk enjoy a smaller risk benefit from renting than households with more exposure do (e.g., one- versus two-earner households), and retirees enjoy a benefit that is smaller still; in fact, for some retirees owning would be safer. Hence, the risk-minimizing life cycle profile involves renting during prime working years and, depending on a household's moving frequency, potentially 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 quantify the effects of owning versus renting housing on the variability of housing costs and net income, and the welfare consequences. This complements and extends previous research on this question in three key ways. First, we utilize over 70 years of data to characterize the risk process over life cycle-relevant horizons. Second, we account for the strong relationships between housing prices and wages. Third, we quantify, using a variety of approaches, the welfare consequences of different housing risk exposures.[3] As discussed above, our central conclusion that for

[2]This correlation means that a working owner's housing costs tend to be higher when income is lower and lower when income is higher, exacerbating income risk. This correlation also 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 tend to be compensated by higher wages.

[3]A related literature considers how wage risk affects optimal or actual housing tenure choices. For example, Campbell and Cocco (2003) and Yao and Zhang (2005) 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. Sinai (2011) notes that correlations between wages and housing prices may be a key determinant of the riskiness of owning vs. renting. We depart from these papers in our

many households renting is safer than owning, even for households with long horizons in a home, contrasts with widely-held views in the research literature, in financial advice, and among households. Our findings highlight important dimensions of heterogeneity, some of which also differ from prevailing views. Although the riskiness of owning versus renting is just one determinant of the optimal housing tenure for any given household, our results suggest that it is a quantitatively important determinant that is widely misunderstood.

Our findings 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. Our back-of-the-envelope calculations suggest that misperceptions about housing risks could increase the lifetime efficiency cost of housing subsidies by about $4,000 per household and reduce welfare by even more.

This paper also contributes new evidence on labor earnings risk, including its determinants and welfare consequences. This complements and extends the large literature on labor earnings risk (see Guvenen et al., 2021, for a recent review). Whereas much of this literature focuses on characterizing individual- or household-level idiosyncratic risk over relatively short time horizons, we focus on market-level risk over longer horizons of up to 70 years. While idiosyncratic risk drives much of the risk in household earnings over short horizons, we find that market-level risk—risk in local labor market conditions—drives much of the risk in a household's earnings over its lifetime. Local labor market shocks can cause differences in wage growth across different locations that are sufficiently large and persistent as to generate substantial risk in lifetime earnings. Our findings also help clarify the welfare and behavioral effects of wage risk by highlighting the key role of housing choices in determining a household's net exposure to wage risk.

2 Theory: Housing Tenure and Net Income Risk

This section uses a simple model to illustrate how wage risk can fundamentally transform the riskiness of owning versus renting housing. Consider a simple model in which a household lives for T periods. In each period $t \in \{1, 2, ...T\}$, it supplies one unit of labor

focus on the welfare cost of different housing exposures, which we investigate by analyzing the evolution of wages, rents, and home prices over life cycle-relevant time periods.

perfectly inelastically, earning wage w_t, and demands one unit of housing, also perfectly inelastically. If the household rents, it pays the going rent R_t in each period. If it owns, it pays P_0 to buy a home immediately before period one and receives P_T when it sells its home after living in it in period T (received after R_T is realized, i.e., the ex dividend price). The present value of non-housing consumption, C, equals the present value of net income, i.e., the present value of income, Y, less the present value of housing costs, H:

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

where r is the household's discount rate.[4]

Consumption risk, measured as the variance of the present value of consumption, can be written[5]

$$Var(C) = Var\left(\sum_{t=1}^{T} \frac{w_t}{(1+r)^t}\right) + \begin{cases} Var\left(\sum_{t=1}^{T} \frac{R_t}{(1+r)^t}\right) & -2Cov\left(\sum_{t=1}^{T} \frac{R_t}{(1+r)^t}, \sum_{t=1}^{T} \frac{w_t}{(1+r)^t}\right) \text{ if rent} \\ \underbrace{Var\left(\frac{P_T}{(1+r)^T}\right)}_{\text{"Direct effect"}} & \underbrace{+2Cov\left(\frac{P_T}{(1+r)^T}, \sum_{t=1}^{T} \frac{w_t}{(1+r)^t}\right)}_{\text{"Portfolio effect"}} \text{ if own.} \end{cases} \tag{1}$$

where the risk comes from risk in wages and housing prices. The riskiness of a given housing exposure depends not only on its riskiness as viewed in isolation, the "direct effect" $Var(H)$, but also on its covariance with the rest of the household's portfolio, the "portfolio effect" $Cov(Y, H)$ (adopting the terminology of Lockwood, 2021). Other things equal, housing costs that covary positively with income are less costly because they reduce the welfare cost of income risk.

In the special case in which housing prices are uncorrelated with wages, the riskiness of owning versus renting depends solely on the "direct effect"—the variance of housing costs. It amounts to a comparison of the variance of the present value of rents during the household's stay to the variance of the present value of the home price at the end of

[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. Much of the literature on the risk properties of owning versus renting housing ignores the risk properties of carrying costs. We exclude them from this analytical section to simplify the exposition, but we allow for them in the quantitative analyses to follow.

[5] XXX- brief note about this simple measure of risk

the household's stay:

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

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, an 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: $V(C_{rent}) - V(C_{own}) \to V(R) > 0$.[6] This logic, due to Sinai and Souleles (2005), supports the prevailing view that for households with long enough horizons in a home, owning is safer than renting.

More generally, the riskiness of different housing positions depends crucially on the covariance between housing prices and wages, the portfolio effect. The portfolio effect can have a major impact on the riskiness of owning versus renting because renters and owners have opposite-signed exposures to housing prices and hence opposite-signed portfolio effects. 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. If housing prices are positively correlated with wages, the portfolio effect reduces the riskiness of renting (since rent risk hedges wage risk) and increases the riskiness of owning (since home price risk exacerbates wage risk). If the correlation is strong enough, the portfolio effect can even make renting safer than full housing cost insurance.[7] In that case, a household that could not otherwise hedge its wage risk benefits from exposure to rent risk and that owning reduces exposure to rent risk is a cost, not a benefit, of owning.[8]

[6] 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 $r = 3\%$.

[7] XXX- math

[8] The portfolio effect also fundamentally alters how changes in horizon affect the relative riskiness of owning versus renting. Lengthening the horizon has an ambiguous effect on the riskiness of owning versus renting if renting provides a valuable hedge against wage risk, since 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.

3 Data and Empirical Approach

Data.— Our main data source is the IPUMS-USA Decennial Census database from 1940 to 2010 (Ruggles et al. 2018). For the purposes of this paper, its main strengths 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 particular is crucial for learning about the risk facing households over life cycle horizons.[9]

We construct market prices from household-level data in 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 local markets. They also have the advantage of being amenable to long-run analyses (Eckert et al. 2018), since county boundaries are stable over time, whereas the boundaries of other common definitions of local labor markets change frequently over time (e.g., MSAs). 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). We drop Alaska and Hawaii because their county population data are not available before they obtained statehood in 1959. This leaves us with 720 out of 744 CZs 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 is between 25 and 60 years old (i.e., the "representative household member"). For individual-level variables such as wage income or age, we use the values of the representative household member.[10] To measure market wages (i.e., "full-time 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

[9] To the best of our knowledge, no other source can match the Census's length of coverage of the key variables of interest, in particular rents. Moreover, the quality of the resulting price measures appears to be quite good. For example, we find that home prices in the Decennial Census track closely those in home transaction data such as Zillow or FHFA in periods where these samples overlap and the volatility of home price growth rates are very similar (aggregate and location-level); see Appendix B.

[10] 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).

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 and quality. Separately for each Census year, we regress log wage income on fixed effects for household age, and we regress log rents and log home prices on fixed effects for home age, size and other home characteristics. Our CZ-level market price measures are thus stripped of differences in quantity and quality composition. Having completed these regressions, we calculate the market-level prices as the median value in a given market-decade.[11] All analyses use household sampling weights across households and population weights when looking across locations. All dollar values are converted to real 2020 dollars using the CPI - All Urban series. Full details of these steps are in Appendix A.

Table 1 summarizes the resulting measures of wages, rents, and home prices. Columns 1–2 focus on the 2010 cross-section. Columns 4–6 consider the standard deviation of location-specific growth rates in these prices over different time horizons. The results reveal three key patterns. First, variation in the growth rates of these prices over time and across locations is large and persistent, including across locations during a given period of time. For instance, the relative growth rate of wages over a 40-year period (roughly the length of a typical career) has a standard deviation of 25%. Over the same time horizon, relative rent and home price growth rates have a standard deviation of 55% and 100% respectively. Moreover, differences in growth rates are not rapidly reversed. The standard deviation of differential growth across locations roughly doubles from 10–40 years and roughly triples from 10–70 years. This is suggestive that risk in these prices might drive significant risk in living standards over the life cycle.

Second, home prices are much more variable than rents. Forty-year differential growth rates of home prices have a standard deviation of 34 log points, versus 24 log points for 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 et al. (2009), Plazzi et al. (2010)).

Third, both rent and home price growth are highly positively correlated with wage growth at all frequencies. The correlation between rent and wage growth is about 70% at all horizons (both total and location-specific growth) and the correlation between home price and wage growth is 40–60%.

[11] Using medians rather than means minimizes the influence of both extreme values and changes over time in top-coding, both of which could 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.

Approach.— Our goal is to understand the risk properties of different housing exposures: the extent to which they loosen or tighten the budget constraint in different states of the world with higher versus lower marginal utility. The conceptual experiment we have in mind is of a household facing exogenous risk in its income and housing costs from risk in the wages, rents, and home prices in its labor and housing markets.[12] 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—and so living standards—of renters and owners?

The key challenge, common to all analyses of risk, is that one only ever observes the single, realized state of the world for each household. We address this issue using the standard approach of modeling risk (i.e., variation across states of the world) using cross-sectional and time-series variation, with a focus on within-location changes over time: differential changes in the key prices in different locations over the same time period. Specifically, we consider a household that initially faces the wage, rent, and home price of a median full-time working household in 2010. The initial wage is $59,000 per year, the median wage among full-time working households in 2010. The initial home value is $190,000 and the initial rent is $12,000, the median home value and rent, respectively, in 2010 among households with that level of earnings.[13] We assume that the possible future realizations of prices for this household are those that arise from letting this initial price vector grow at the observed location-specific growth paths in each of the 720 locations, where the growth paths in one location are taken to be the realized growth paths in a corresponding state of the world. The resulting risk process isolates the across-location heterogeneity in the growth rates of wages, rents, and home prices over time. It is a highly flexible, non-parametric model of the underlying data generating process. This risk process underlies our baseline analyses in Sections 4 and 5, though we test robustness to other risk processes as well.

[12] In reality, households can re-optimize in response to price changes, which makes risk in prices different from risk in net income (income net of housing costs). 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. 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.

[13] The resulting rent and home value differ from the median rent and home value because the median renter earns less than the median full-time working household and the median owner earns more.

4 Descriptive Analysis of Housing Risk

Fact 1: Volatility in housing costs is minimized by owning if the horizon is long, and by renting otherwise.

Figure 1 plots the standard deviation of the present value of housing costs for owners versus renters as a function of the time horizon in the same home or location. For renters, housing cost volatility increases roughly linearly for short and medium horizons from a low base before eventually decreasing due to the increasing effect of discounting when the tenure horizon becomes very long. For owners, it starts at a high base and decreases in the horizon to a level below that of renters after 18 years. As a result, volatility in housing costs is minimized by owning if the horizon is long and by renting otherwise.[14]

Implications.— For a household that faced housing risk alone or, more generally, whose non-housing portfolio as a whole was statistically independent of housing prices, renting would be safer than owning over short and medium time horizons in a home and owning would be safer than renting over long time horizons in a home. More generally, risk in housing costs is significant, and a household's exposure to that risk depends crucially on whether it owns or rents its home. So the choice of whether to own or rent is a major determinant of a household's overall risk exposure.

Fact 2: Rent growth and home price growth are strongly positively correlated with wage growth over all horizons.

The most important risk for many households is wage risk. Figure 2 shows the joint distributions across locations of rent growth and wage growth (left panels) and home price growth and wage growth (right panels). Associated heterogeneity and robustness analyses are in Appendix B. For all horizons from 10–70 years and all specifications, rent growth and home price growth are strongly positively correlated with wage growth. Ten percent faster than average wage growth over a decade is associated with 13% excess rent growth and 16% excess home price growth. Significant shares of the variation in rents and home prices are associated with variation in wages, with R^2 statistics ranging from 40–50% for rents and 10–25% for home prices (depending on the horizon).[15]

[14]The contrasting effects of changes in the horizon for owners versus renters are driven by the fact that a renter's rent risk accumulates over time, as additional draws of risky rent are realized, whereas an owner's home price risk decreases in the horizon due to discounting overpowering the increasing dispersion in the non-discounted level of the home price over time.

[15]In Appendix B we provide evidence that suggests a fairly strong positive relationship between national-level changes in wages and rents as well, though this is necessarily less certain given that only a

Implications.— Renting hedges wage risk, and owning exacerbates wage risk. So for households exposed to wage risk, the riskiness of a given housing exposure cannot be determined in isolation—based on the volatility of housing costs alone—without regard to its interaction with wage risk. The strong relationship between housing prices and wages means that a household's exposure to wage risk is a crucial determinant of the contribution of a given housing exposure to the overall risk of its portfolio as a whole. That rents and wages covary strongly positively means that for working households, wages and housing costs covary positively if they rent and negatively if they own. In other words, for working households, a renter's housing costs tend to be lower when wages are lower and higher when wages are higher—which mitigates wage shocks—whereas an owner's housing costs tend to be higher when wages are lower and lower when wages are higher—which exacerbates wage shocks. As a result, renting is less risky than it would appear to be from its contribution to risk in the cost-of-living alone, since it hedges wage risk, and owning is more risky than it would appear to be from its contribution to risk in the cost-of-living alone, since it exacerbates wage risk.

Fact 3: For working households, renting reduces net income risk and owning increases it, regardless of the time horizon in the home.

Figure 3 plots the standard deviation of the present value of net income (income minus housing costs) of a typical working household as a function of the time horizon in the home under three housing scenarios: own, rent, and hold full housing cost insurance (so face only wage risk). The figure also shows the net income risk of a household with no exposure to wage risk (e.g., retirees) if they own or rent, reproducing Fact 1 of Figure 1. Regardless of the time horizon in the home, the standard deviation of a working household's net income is minimized by renting, is greater with full housing cost insurance, and is maximized by owning.

How could a renter's exposure to the substantial risk in rents *reduce* their net income risk even relative to full housing cost insurance? Recall from Section 2 that exposure to housing risk can decrease net income risk if the hedge of wage risk dominates the direct effect of housing risk itself. This occurs if the covariance between housing costs and income is sufficiently large relative to the variance of housing costs: $2Cov(Y, H) > Var(H)$, i.e., if the slope of the regression of income on housing costs exceeds one half. We find that for a wide variety of specifications, the slope significantly exceeds one half for renters (e.g., $\widehat{\beta}_{Y|R} = 1.82$ over a 70-year tenure horizon) and is significantly negative

single national time series is observed. We also find similar relationships across countries over long time periods, though the quality of the data is lower and these findings can be sensitive to specifications.

for owners ($\widehat{\beta}_{Y|O} = -2.01$). Hence, net income and housing costs covary positively for renters and negatively for owners. In other words, renters' housing costs tend to be lower when net income is lower and higher when net income is higher—which mitigates net income shocks—whereas owners' housing costs tend to be higher when net income is lower and lower when net income is higher—which exacerbates net income shocks.

Decomposing the slope for renters into $\frac{Cov(Y,R)}{Var(R)} = Corr(Y,H) \times SD(Y)/SD(H)$, the hedge of wage risk dominates the direct effect of rents being risky because (i) a renter's housing costs are a fairly good hedge against wage risk ($Corr(Y,R) \approx 0.75$), and (ii) variation in wages far exceeds that in a renter's housing costs ($SD(Y)/SD(R) \approx 2.4$).[16] In second-best contexts in which people face uninsured risks beyond housing costs, greater exposure to housing cost risk does not imply greater exposure to risk on net. Given the substantial risk in wages, the hedge against wage risk provided by renting significantly reduces the net income risk of working households (moderate reduction in a large wedge), more than its contribution to cost-of-living risk increases it (moderate increase in a smaller wedge).

This connects the three main findings of our analysis of housing risk. Renting reduces net income risk relative to full housing cost insurance (Fact 3) because its valuable hedge against wage risk (Fact 2) dominates its moderate contribution to cost-of-living risk (Fact 1), which in turn is largely because wage risk is greater than rent risk. Owning increases net income risk relative to full housing cost insurance because of reinforcing effects of increasing cost-of-living risk, especially if the horizon is not long (Fact 1), and exacerbating wage risk (Fact 2). As a result, the net income of working households is much more volatile if they own than if they rent.

For a different way of understanding the impact of owning versus renting on net income risk, Figure 4 shows graphically how owning and renting affect the extent to which shocks to income translate into shocks to net income. The strong positive correlations between wages and housing prices create a positive correlation between the earnings and housing costs of renters of 0.75 and a negative correlation between the earnings and housing costs of owners of -0.30. As a result, a renter's exposure reduces the extent to which income shocks translate into net income shocks, reducing the slope of net income with respect to income to 0.43, whereas an owner's exposure increases the extent to which income shocks translate into net income shocks, increasing the slope of net income with respect to income to 1.02.

Implications.— Regardless of the time horizon in the home, renting reduces a working

[16]The standard deviation of the present value of rents (wages) is about \$4,000 (\$9,000) after one year, \$48,000 (\$78,000) after 15 years, and \$86,000 (\$203,000) after 40 years.

household's net income risk and owning increases it. The variance of net income is lowest for renters, somewhat greater for households with full housing cost insurance, somewhat greater still for owners with long horizons, and even greater still for owners with shorter horizons. Owning is riskier than it has been thought to be because not only does it exacerbate wage risk, it also eliminates exposure to rent risk and so the valuable hedge against wage risk that rent risk provides. That is, for working households, exposure to rent risk is not costly, as has been thought, but valuable.

Robustness.— Facts 1–3 are highly robust, including to using different sample periods (e.g., excluding the Great Recession, the recent housing boom and bust cycle, or the post-war economic boom), different market definitions (e.g., counties or MSAs instead of CZs), and different measures of income and housing costs (e.g., total family income or gross instead of contract rents). They are robust to controlling for population changes and dropping households that migrated between CZs. They also hold across the distribution of wages and housing costs (e.g., at the 25^{th} or 75^{th} 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 shorter sample 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 across many CZs) are not associated with location-specific wage changes. Full details of these robustness checks are in Appendix B.

5 The Risk Cost of Owning Versus Renting

5.1 Estimates of the risk cost of owning versus renting

In this section, we go beyond the preceding summaries of the raw data to quantify the ex ante welfare costs of the risk aspect of owning versus renting.

Approach.— We combine the evidence on the risk in wages, rents, and home prices from Section 4 with various assumptions about the determination of consumption and the utility function to estimate the risk cost of different housing exposures.

Risk.— Our baseline approach to modeling the risk that the household faces is that described in Section 3. We consider a household that faces risk in its income and housing costs arising from risk in wages, rents, and home prices. Risk in 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. For simplicity and to be conservative with respect to our key conclusions, we assume that owners' housing costs are perfectly smoothed over the life cycle, equal to the annuitized value of lifetime housing costs in each period. This gives a lower bound on the cost of an owner's housing risk exposure.

Consumption.— We use the following three models of consumption:

- Hand-to-mouth (HTM): $c_{\omega,t} = y_{\omega,t} - h_{\omega,t}\ \forall \omega, t$
- Perfect within-lifetime consumption smoothing (PS): $c_{\omega,t} = \bar{c}_{\omega} \equiv \frac{PV_{\omega}(Y) - PV_{\omega}(H)}{\sum_{i=1}^{T}(1+r)^{-i}}\ \forall \omega, t$
- Life-cycle model (LC): $c_{\omega,t} = c^{\star}_{\omega,t}\ \forall \omega, t$, where $c^{\star}_{\omega,t}$ is the optimal consumption in a life-cycle model

where ω is the state of the world and t is the time period. The HTM and PS models share the same key strengths: transparency and flexibility in accommodating any possible model of the risk process. The HTM model features no consumption smoothing over time. As such, it likely is a poor model of behavior in the face of short-term, transitory shocks but a better model of behavior in the face of the kind of longer-term, more persistent risks that matter most for lifetime welfare. It can be a useful benchmark in part because it likely overstates the welfare cost of fluctuations in net income over time, so the associated risk costs can be viewed as an upper bound. The PS model is the opposite extreme in that, by minimizing the welfare cost of a given amount of risk in the present value of net income, it should put a lower bound on the risk cost or benefit of different housing exposures. The LC model has the benefit of allowing for partial consumption smoothing within lifetime states of the world based on information revealed up to that point, at the cost of requiring many strong assumptions for implementation, including about the underlying risk process. It also accounts for borrowing constraints that households might face. See Appendix C for details. In the main text, we focus mostly on the PS model for simplicity and to be conservative, as it tends to understate the risk benefit of renting relative to owning (by understating both the risk cost of owning and the risk benefit of renting).

Utility.— We assume that the flow utility function is constant relative risk aversion. The baseline coefficient of relative risk aversion is two, and we test robustness to alternatives.

Risk cost calculations.— Our main measure of risk cost is the equivalent variation of going from the housing exposure of interest to a counterfactual alternative in which housing costs are fully insured at actuarially fair prices. In particular, we report the present value of the uniform increment to income in all time periods under the housing exposure

of interest that makes the household as well off ex ante (same expected utility at the beginning of the life cycle) as it would be if housing costs were fully insured at actuarially fair prices. The full-insurance counterfactual for renters involves the household paying at each date the expected value of rents at that same date. So to the extent that the expected value of rents changes over time, the full-insurance counterfactual for renters features the same changes in the certain rent over time. This means that the only thing that changes between the renter's status quo exposure and their counterfactual full-insurance exposure is risk, not changes in expected rents over time. For owners, at each date the household pays the same annuitized value of the expected present value of lifetime net housing costs. Hence, owners' housing costs are perfectly smoothed over the life cycle in both the status quo and full-insurance counterfactual. We make this assumption, which understates the riskiness of owning, in order to be conservative with respect to our key conclusion that owning tends to be riskier than renting. For both renters and owners, the expected present value of housing costs is the same under the risky and certain housing cost counterfactuals, so the equivalent variation of going from risky to certain housing costs is a measure of the pure risk cost of the risky housing exposure.

For each set of assumptions, we calculate the lifetime risk costs of the exposures to housing risk, of households that move at various frequencies, from the following life cycle profiles of housing tenure status: always rent, always own, rent during working years and own during retirement, and rent during working years and hold full, actuarially fair housing cost insurance during retirement (which is similar in terms of risk properties to taking out a certain type of reverse mortgage at the beginning of retirement). Our main analysis assumes that moves are exogenous in terms of both timing and destination. For timing, we consider many different moving frequencies. For destination, we draw the destination randomly from the set of locations, with the probability of each location proportional to its population.

Results.— Figure 5 plots, for the PS consumption model, the risk costs of owning and renting, both relative to full housing cost insurance, as a function of the frequency of moves over the life cycle. Owning has a positive risk cost regardless of moving frequency. The risk cost of owning first increases in the time between moves as risk accumulates, reaching a maximum of about \$4,000 for moving frequencies of one move every 20–40 years, before decreasing once time discounting starts to dominate risk accumulation. By contrast, renting has a negative risk cost (positive risk benefit) relative to full housing cost insurance regardless of moving frequency. This risk benefit increases roughly linearly in the time between moves as risk accumulates, reaching a maximum risk benefit of over \$8,000 if the household never moves. As a result, regardless of the moving frequency,

renting provides a substantial risk benefit over owning.

These conclusions about the relative riskiness of different life cycle housing tenure profiles are highly robust to a wide range of changes in the risk process, the model of consumption, and other changes to the model (associated tables and figures coming soon).

5.2 Key driver: Strong correlations between housing prices and wages from local labor demand shocks

Proximate driver: Correlated wage risk.— Figure 6 plots, for the PS consumption model, the lifetime risk costs of owning and renting for a household that never moves as a function of that household's exposure to wage risk. The measure of exposure to wage risk is the "wage risk beta" β_w, where the wage in state-time (ω, t) is $w_{\omega,t}(\beta_w) = \beta_w w_{\omega,t} + (1-\beta_w)E(w_{\omega,t}|t)$. So a wage risk beta of two is not like a two-earner household but a one-earner household whose wages are the same as a typical household on average but twice as volatile around the same mean. A beta of zero means the household faces no wage risk. A beta of one is the baseline risk model.[17]

The results reveal the crucial role of the interaction with wage risk in driving the key conclusions. The interaction with wage risk transforms a renter's exposure to risk in housing costs from being quite costly (welfare cost of \$7,300 when $\beta_w = 0$) to quite valuable (welfare benefit of \$8,400 when $\beta_w = 1$). This is because renting provides a highly valuable hedge against earnings risk (worth \$15,700) that dominates the risk cost of its increase in cost-of-living risk (cost of \$7,300). This is in keeping with Fact 3 in the descriptive analysis that renting has a valuable portfolio effect from hedging income risk that dominates its partial effect cost from increasing cost-of-living risk and so reduces net risk exposure even relative to full insurance against housing cost risk. For an owner, by contrast, the interaction with wage risk increases the welfare cost of their exposure to risk in housing costs. Even for a household that never moves, the exacerbation of wage risk from the fact that an owner's capital gain from housing is strongly positively correlated with wages imposes a welfare cost of \$400 on top of the \$80 welfare cost from the contribution of housing cost risk to cost-of-living risk.

[17] Of course, differences in wage risk beta lead to differences not only in the extent to which wages covary with housing costs but also in the overall level of wage volatility. An alternative would be to hold constant the level of wage volatility while varying only the correlation with housing costs. In practice, we find that the risk cost of different housing exposures are virtually identical if we eliminate wage volatility entirely ($\beta_w = 0$) or if we hold fixed the marginal distribution of wages at its baseline level but eliminate the correlation with housing costs, so we use the simpler notion of wage beta.

In summary, the key proximate driver of the conclusion that for working households renting is safer than owning regardless of the time horizon in the home is the strong positive correlations between home prices and wages. Because of these correlations, renting hedges wage risk and owning exacerbates it. Were it not for wage risk, a renter's housing exposure would have a straightforward effect on consumption risk relative to full housing cost insurance: It would increase consumption risk due to its contribution to cost-of-living risk. But because of the strong relationship between rents and wages, a renter's housing exposure effectively bundles two assets: one that pays off more when rents are lower, which increases the household's cost-of-living risk, and one that pays off more when wages are lower, which hedges wage risk. The overall effect on risk exposure relative to full housing cost insurance therefore reflects opposing pro- and anti-insurance effects. In principle, either could dominate. In practice, the insurance of wage risk dominates the increase in cost-of-living risk, so renting decreases risk exposure on net even relative to full housing cost insurance. As a result, for working households renting is significantly safer than owning regardless of the time horizon in the home.

Ultimate driver: Local labor demand shocks.— Why are location-specific changes in housing prices and wages so strongly positively correlated? Although this strong relationship likely reflects multiple factors, a key one is likely to be the combination of (i) significant location-specific shocks to labor demand, (ii) relatively inelastic supply of housing and other location-specific capital in many locations, and (iii) spatial equilibrium forces toward different locations being similarly attractive to potential residents.

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 et al., 2013; Giannone, 2018; Howard and Liebersohn, 2018; Hornbeck and Moretti, 2019).

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 importance in explaining the relationship between location-specific changes in wages and housing prices, by shifting the demand for housing. The strength of the association between location-specific changes in wages and housing prices suggests that wage changes shift the demand for housing substantially 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 (Glaeser and Gyourko, 2018).

The apparently large impact of wage changes on the demand for housing likely primarily reflects a combination of income effects (wage growth enriches working households, which increases their demand for normal goods, including housing) and, especially, spatial equilibrium-type responses to location-specific wage changes. Locations with high wages tend to have high housing costs, and the positive relationship between wages and housing costs is even stronger when isolating location-specific growth in these prices relative to the national average. This is consistent with spatial equilibrium forces ensuring that, other things equal, places that experience faster growth in wages experience faster growth in the prices of non-tradables like housing. The strong positive correlation between the growth of a location's wages and rents relative to the national average provides additional support for the conclusion that local labor demand risk is the major driver of location-specific risk, since risk in local consumer amenities, by contrast, would tend to lead to rents covarying *negatively* with wages (Roback, 1982).

6 Implications

Is renting better than owning? *If* the key difference between owning and renting were the effect on risk exposure, then renting would be significantly better than owning for most households. A rough, back-of-the-envelope calculation based on our estimates of the risk cost of owning and the risk benefit of renting suggests that for a typical household, the lifetime risk benefit of renting rather than owning is on the order of $10,000. But risk is not the only difference between owning and renting. Other differences include any preference people might have for owning their own home, moral hazard and other transaction costs between renters and landlords (including the risk of non-renewal of a rental contract), and considerations related to mortgages and owned housing as a store of wealth and saving commitment. Given the many differences between owning and renting, that renting is safer than owning just means that one important component of the overall welfare effect of owner-occupied housing is not the benefit it has been thought but a cost.

How generalizable are the results? The analysis highlights the importance of two factors that likely vary significantly across households both within and across different economies: exposure to wage risk and time horizon in a home or location. The smaller the exposure to wage risk, the smaller the risk benefit from renting and the risk cost from owning. Still, our results suggest that, other things equal, exposure to wage risk

would have to be significantly smaller than that of a typical working-age household in the U.S. to make owning safer than renting, and even then its risk advantage would be small. The key factor is that wage risk is a major risk—for many households the biggest risk they face—so even a modest hedge against it can be quite valuable and a modest exacerbation of it can be quite costly.

The time horizon in a home or location is also an important determinant of the risk cost of owning versus renting, though it has a more subtle effect than previously thought. In addition to the well-known effect due to time discounting that a longer horizon eventually reduces the sale price risk of owning, there is an opposing effect due to the interaction with wage risk that a longer horizon increases the value of the wage risk hedge from renting. Our results suggest that the net effect of these and other forces is such that renting is safer than owning for working households regardless of the horizon and for retired households with horizons of up to roughly 30 years.

What are the 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 (e.g., Glaeser and Gyourko, 2018). Such restrictions, by reducing the elasticity of housing supply, increase the extent to which labor demand shocks capitalize into wages and housing costs, which increases the risk in earnings and housing costs. They also increase 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, and that these misperceptions affect choices about whether to own or rent. 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. In a nationally representative survey of more than 50,000 households, Adelino et al. (2018) find that 71% of U.S. households believe that housing is a "safe" investment, significantly higher than the 55% who believe the same about government and corporate bonds. Moreover, Adelino et al. (2018) find that both current housing tenure status and future intentions to buy versus rent are strongly correlated with perceptions of house price risk. In light of our finding that for many households owning is riskier than renting, these survey results suggest that many people are underestimating the riskiness of owning relative to renting and that these misperceptions affect choices about whether to own or rent. If so, such misperceptions 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.

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 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.[18] 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.)

Internality wedge σ_{int}.— 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 household that never moved would underestimate the *lifetime* (risk) cost of owning relative to renting by more than \$16,000: $\sigma_{int} = \$16,000$.

[18] This is the marginal excess burden if there are no spillovers to other distorted markets and if the difference between the social marginal cost and social marginal benefit of owner-occupied housing is linear in the quantity of owner-occupied housing (as would be the case if both of those schedules are linear). If these schedules are non-linear, this is a second-order approximation to the marginal excess burden.

Net subsidy σ_{sub}.— A variety of policies subsidize owner-occupied relative to rental housing, and the effective subsidy 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. As a rough approximation, we assume that the net subsidy for a typical household is the product of its marginal income tax rate and its (possibly implicit) rent, $\sigma_{sub} = \tau \times R$. This is a rough approximation to what is likely the most important driver of the net subsidy for owner-occupied housing: that the implicit rent of owners is untaxed, whereas the explicit rental income of landlords is taxed.[19] We assume a benchmark marginal tax rate on income τ of 25 percent and an annual rent R of \$14,700, which yields an annual subsidy on owner-occupied housing of roughly \$3,500 per year.[20] We assume that the net subsidy *per lifetime* of living in owner-occupied housing σ_{sub} (to match the units of σ_{int}) is \$90,000, the value of a 50-year annuity of \$3,500 per year with a 3 percent annual discount rate.

Quantity response ΔQ_{sub}.— For our benchmark quantity response, we assume that the subsidy increases the fraction of its lifetime that a typical household owns by 25 percentage points $\Delta Q_{sub} = 0.25$, roughly the difference in the homeownership rate between Switzerland, which has small homeowner subsidies, and 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. We also consider quantity responses half as large: $\Delta Q_{sub} = 0.125$.

Results.— Table 2 shows the results. The estimates suggest that misperceptions increase the efficiency cost of ownership subsidies by about one-third, roughly from \$11,000 to \$15,000 per household over its lifetime.[21] Misperceptions themselves have a welfare cost of over \$4,000 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-

[19]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. 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. While it is impossible to capture the full richness of the tax treatment of housing in a simple calculation, our aim is to provide rough, back-of-the-envelope estimates of the likely magnitude for a typical household.

[20]This annual rent comes from applying the average net annual rental yield on housing of 4.2 percent (Demers and Eisfeldt, 2022) to the median home price in 2020 of \$350,000 (based on data from FRED). This estimate of the annual subsidy is broadly in keeping with those of Gyourko and Sinai (2003), who estimate per-owner tax subsidies on the order of \$1,000–\$6,000.

[21]While the absolute size of the efficiency cost depends on (is proportional to) the quantity response ΔQ_{sub} (the most uncertain ingredient), the implicit subsidies alone determine the proportional increase due to misperceptions: $MEB_{sub}/MEB_{sub|\sigma_{int}=0} = 1 + 2\sigma_{int}/\sigma_{sub}$.

optimal, and indirect, by magnifying the efficiency cost of policies that subsidize owner-occupied relative to rental housing beyond the associated externalities.

7 Conclusion

For many U.S. households, renting is safer than owning, often significantly so. The key driver is wage risk. Because of the strong positive correlations between changes in home prices and changes in wages, renting hedges wage risk and owning exacerbates it. This is an important component of the overall welfare effect of owning versus renting and associated policies.

Our findings have 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 about the performance of different models. 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.

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.

Baker, Scott R., Pawel Janas, and Lorenz Kueng, "Correlation in State and Local Tax Changes," *Available at SSRN 3648513*, 2020.

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

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

Campbell, John Y. and João F. Cocco, "Household Risk Management and Optimal Mortgage Choice," *Quarterly Journal of Economics*, 2003, *118* (4), 1449–1494.

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.

Demers, Andrew and Andrea L. Eisfeldt, "Total returns to single-family rentals," *Real Estate Economics*, 2022, *50* (1), 7–32.

Eckert, Fabian, Andres Gvirtz, and Michael Peters, "A Consistent County-Level Crosswalk for US Spatial Data since 1790," Technical Report, Working Paper 2018.

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.

Guvenen, Fatih, Fatih Karahan, Serdar Ozkan, and Jae Song, "What Do Data on Millions of US Workers Reveal About Life-Cycle Earnings Risk?," *Econometrica*, 2021, *89* (5), 2303–2339.

Gyourko, Joseph and Todd Sinai, "The Spatial Distribution of Housing-Related Ordinary Income Tax Benefits," *Real Estate Economics*, 2003, *31* (4), 527–575.

Hornbeck, Richard and Enrico Moretti, "Who Benefits From Productivity Growth? Direct and Indirect Effects of Local TFP Growth on Wages, Rents, and Inequality," *Working Paper*, 2019.

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

Lockwood, Lee M, "Health Insurance Increases Financial Risk," 2021. University of Virginia, mimeo.

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

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.

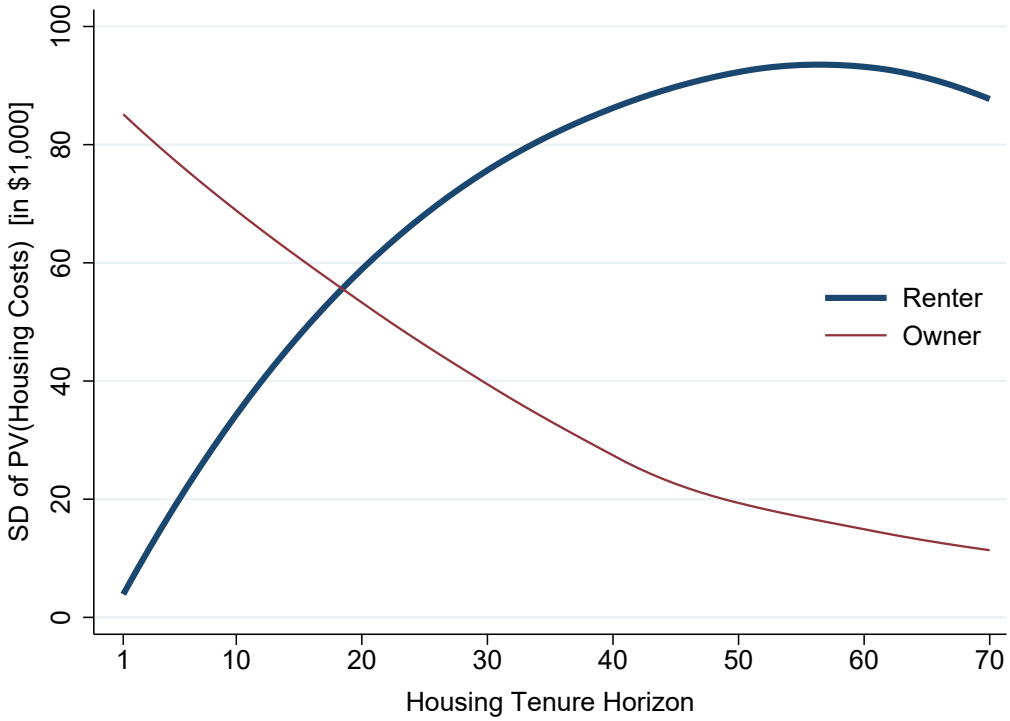
Ruggles, Steven, Sarah Flood, Ronald Goeken, Josiah Grover, Erin Meyer, Jose Pacas, and Matthew Sobek, *IPUMS USA: Version 8.0 [dataset]*, Minneapolis, MN: IPUMS, 2018.

Sinai, Todd, "Understanding and Mitigating Rental Risk," *Cityscape: A Journal of Policy Development and Research*, 2011, *13* (2).

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

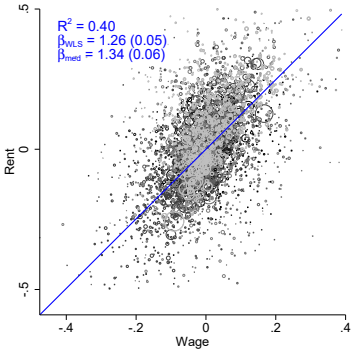
Yao, Rui and Harold H. Zhang, "Optimal consumption and portfolio choices with risky housing and borrowing constraints," *Review of Financial Studies*, 2005, *18*, 197239.

Figure 1: Risk in housing costs for renters and owners by tenure horizon

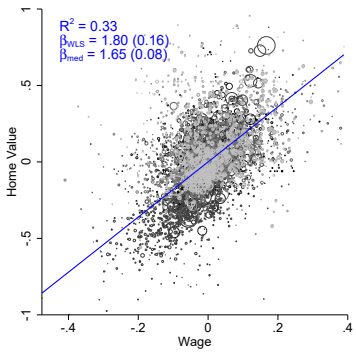


Notes. Standard deviation of the present value of housing costs by tenure horizion T across locations c. Rents and home prices grow at location-level growth rates relative to population-weighted national growth rates. Renters' present value of housing costs in location-year (c,t) are $R_{c,t} = \sum_{s=1}^{T} \frac{R_{c,t+s}}{(1+r)^s}$; the owners' housing costs are $O_{c,t} = P_t - \frac{P_{c,t+T}}{(1+r)^T}$. Initial purchase price is known and does not contribute to the standard deviation of owners' costs. The number of observations ranges from 50,400 location-periods for a one-year horizon (70 one-year periods in each of 720 locations) to 720 location-periods for a 70-year horizon. The discount rate r is set to 3%. All variables are residualized as described in the text before calculating present values. All values are in dollars of 2020. Source: Decennial Census and American Community Survey (ACS).

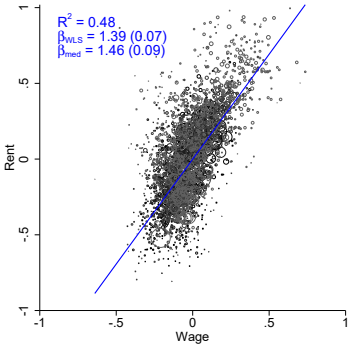
Figure 2: Relationships between Wages, Rents, and Home Prices



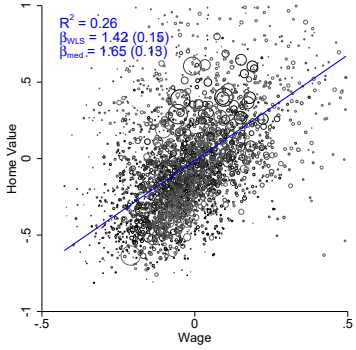
(a) rent/wage: 10-year growth



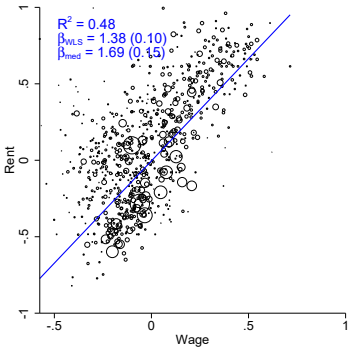
(b) price/wage: 10-year growth



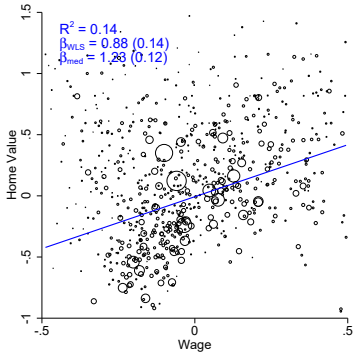
(c) rent/wage: 40-year growth



(d) price/wage: 40-year growth



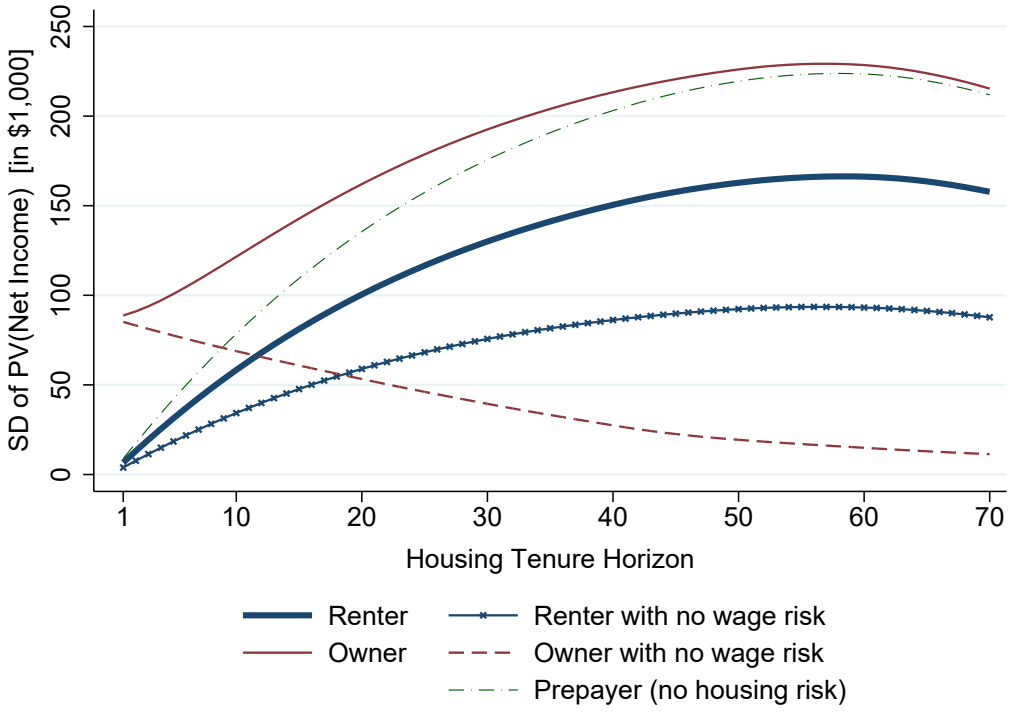
(e) rent/wage: 70-year growth



(f) price/wage: 70-year growth

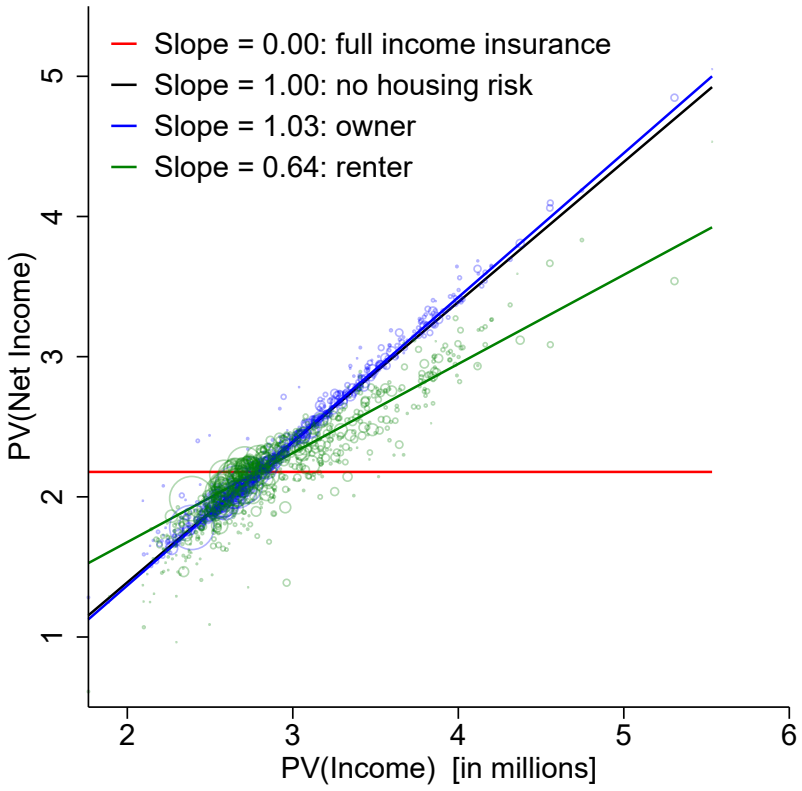
Notes. Growth rates of median full-time wages, rents and home values by Commuting Zone (CZ) relative to national level growth rates (i.e., demeaned). Circle areas represent CZ population size relative to total population in each decade. Observations in darker shades are from more recent decades. β_{WLS} and R^2 are obtained from regressing growth rates in housing costs on growth rates in wages, weighted by CZ population. β_{med} is the corresponding median regression (also population-weighted) to assess robustness of the least square coefficient to outliers. Source: Decennial Census and American Community Survey (ACS).

Figure 3: Risk in net income for renters and owners by tenure horizon



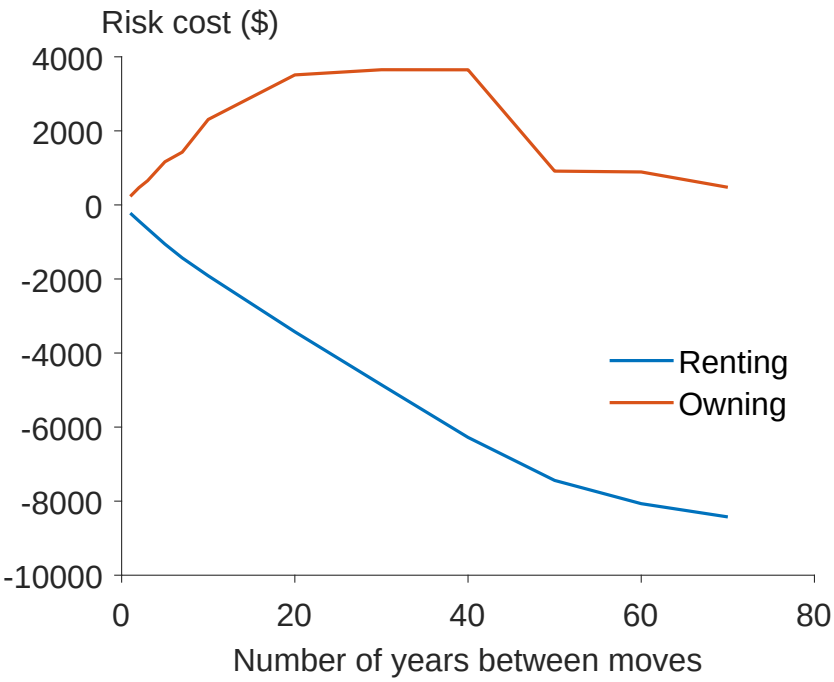
Notes. See Figure 1 for details. Net income is full-time wage income net of housing costs. Renters with no wage risk receive the expected wage in each year, i.e., the average wage across all 720 locations in a year. The prepayer pays the expected rent in each year, i.e., the average wage across all 720 locations in a year. Source: Decennial Census and American Community Survey (ACS).

Figure 4: Insurance of Lifetime Income Risk



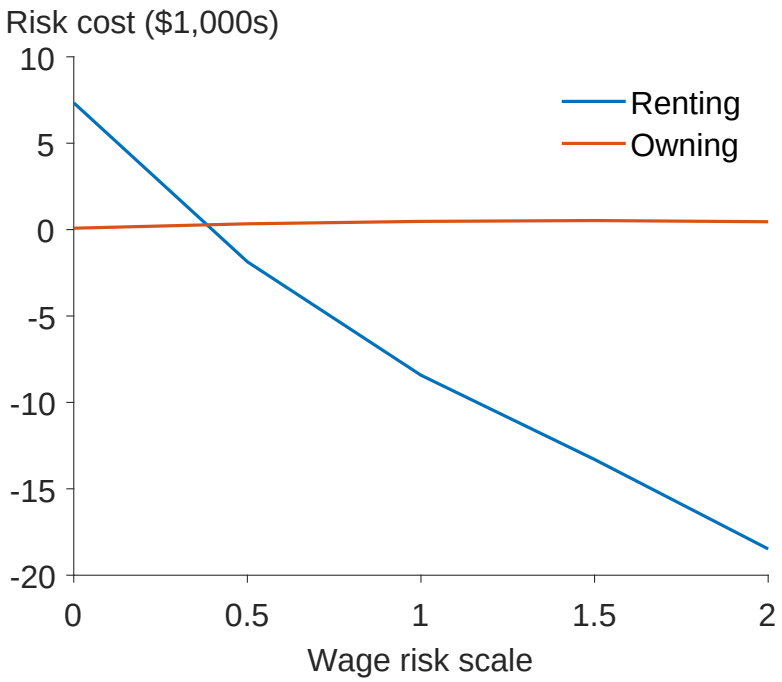
Notes. Lifetime present values of median full-time wage income of a household that stays in the same location for seventy years are shown on the horizontal axis. Corresponding lifetime present values of non-housing consumption (i.e., net income) when renting, prepaying, or owning housing are on the vertical axis. An observation is a CZ-present value using 720 CZs from 1940 to 2010. All variables are residualized as described in the text before calculating present values. Circle areas represent CZ population size relative to total population in each decade. Source: Decennial Census and ACS.

Figure 5: Lifetime Risk Cost as Function of Time Between Moves



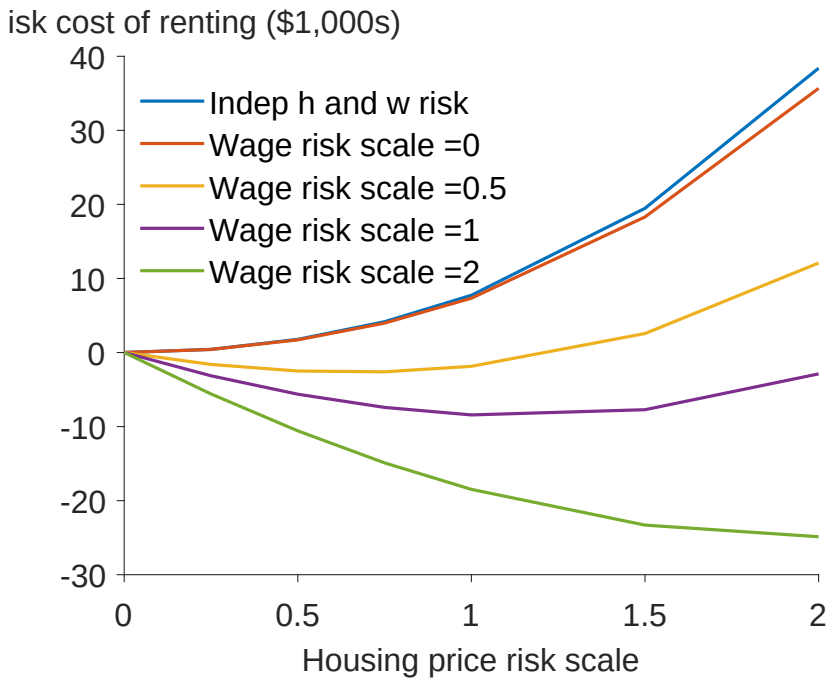
Notes. Lifetime risk costs of owning and renting housing, respectively, as a function of the number of years between moves during the life cycle. See Section 5 for details.

Figure 6: Lifetime Risk Cost as Function of Wage Risk Exposure



Notes. Lifetime risk costs of owning and renting housing, respectively, as a function of a household that never moves as a function of that household's exposure to wage risk. The measure of exposure to wage risk, the "wage risk scale" α, sets the wage in state-time (ω, t) to $Y_{\omega,t}(\alpha_Y) = \alpha_Y Y_{\omega,t} + (1 - \alpha_Y)E(Y_{\omega,t}|t)$. So a wage risk scale of two is not like a two-earner household but rather like a one-earner household with a "wage beta" of two, i.e., a household whose wages are the same as a typical household on average but twice as volatile around the same mean. A wage risk scale of zero means the household faces no wage risk. A scale of one is the baseline risk model. See Section 5 for details.

Figure 7: Lifetime Risk Cost as Function of Wage Risk Exposure



Notes. Lifetime risk costs of owning and renting housing, respectively, as a function of a household that never moves as a function of that household's exposure to wage risk. The measure of exposure to rent risk, the "rent risk scale" α, sets the rent in state-time (ω, t) to $R_{\omega,t}(\alpha_R) = \alpha R_{\omega,t} + (1-\alpha_R)E(R_{\omega,t}|t)$. So a rent risk scale of two is like renting a home twice the size (or quality) of that of the median income household in that location, i.e., a household whose rents are the same as a typical household on average but twice as volatile around the same mean. A rent risk scale of zero means the household faces no rent risk. A scale of one is the baseline risk model. See Section 5 for details.

Table 1 – Wages, rents and home prices

	2010 cross-section:			StDev of location-specific growth rates (log changes):		
	Mean	SD	Corr(x,Y_t)	10 years	40 years	70 years
Wage, $Y_{c,t}$	58,707	10,160	100.0%	8.0%	25.1%	46.4%
				(6.6%)	(12.3%)	(18.5%)
Rent, $R_{c,t}$	10,807	3,346	82.5%	15.9%	54.6%	115.3%
				(13.0%)	(24.8%)	(36.8%)
Home value, $P_{c,t}$	242,894	128,618	81.9%	25.8%	100.7%	165.2%
				(20.6%)	(34.5%)	(43.4%)

Notes: All calculations use population weights for the 720 commuting zones (CZs). Annual full-time wages, annual rents and home values are CZ medians and residualized as described in the text. Columns 1 to 3 use the 2010 cross-section with values expressed in dollars of 2020. Columns 4 to 6 pools data from 1940 to 2010 so that the standard deviation of 10-year growth rates (log changes) is based on 5,040 CZ-year observations respectively 2,880 for 40-year growth rates (log changes) and 720 for 70-year growth rates (log changes). Growth rates (and log changes) are location-level changes relative to population-weighted national changes. All variables are residualized as described in the text. Source: Decennial Census and ACS.

Table 2: Efficiency costs of subsidies and misperceptions

Panel A: Ingredients

Implicit subsidies ($s per lifetime)		Quantity response (in lifetimes)		
			β^{low}	β^{main}
Misperceptions, σ_{int}	16,000			
Net subsidy, σ_{sub}	90,000	$\Delta Q_{sub} = \beta\sigma_{sub}$	0.13	0.25
		$\Delta Q_{int} = \beta\sigma_{int}$	0.02	0.04

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

	With misperceptions		Without misperceptions		Effect of misperceptions	
Formula	β^{low}	β^{main}	β^{low}	β^{main}	β^{low}	β^{main}
$\left(\sigma_{int} + \frac{\sigma_{sub}}{2}\right) \Delta Q_{sub}$	7,625	15,250	5,625	11,250	2,000	4,000

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

	With subsidies		Without subsidies		Effect of subsidies	
Formula	β^{low}	β^{main}	β^{low}	β^{main}	β^{low}	β^{main}
$\left(\sigma_{sub} + \frac{\sigma_{int}}{2}\right) \Delta Q_{int}$	2,178	4,356	178	356	2,000	4,000

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^{main} (β^{L}) is the slope of demand for owner-occupied housing with respective to the subsidy rate assuming the baseline (low) 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 Data

A.1 Location-Level IPUMS Data

We use a 5% random sample, stratified by county, of the IPUMS 1940 100% "full count data" (to speed up the computations), 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 2012 5-year ACS 5% sample, all downloaded from https://usa.ipums.org/usa (Ruggles et al. 2018). The 1950 1% and the 2010 10% sample lack data on rent, income, and home values. We therefore use the 2012 5-year ACS for 2010 – which is centered at 2010, covering years 2008-2012 – and we interpolate values in 1950 log-linearly between 1940 and 1960. Our results are robust to dropping 1950 from our sample or to using the shorter period from 1960-2010. When using population weights, we use accurate CZ- or county-level populations obtained from the full Census counts instead of estimates based on the IPUMS samples.

A.1.1 Defining Local Labor Markets for 1940-2010

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 local 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 change frequently over time (e.g., MSAs).

The most detailed geolocation available in the microdata is a Public Use Microdata Area (PUMA) or a County Group, which maps either to a single county or a group of coun-

ties, 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 or where PUMAS are larger than counties and potentially cross CZs. 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. These crosswalk files and code are available on the authors' website.

Although county boundaries are fairly stable during our sample period, some changes do occur. To take these into account, we use the crosswalk files from Eckert et al. (2018) which maps counties in each decade to stable county identifiers based on county areas in 1990. We then assign each county to a unique CZ using the definition of CZs in 1990.

We drop Alaska and Hawaii because their county population data are not available before they obtained statehood in 1959. We also drop two small CZs with missing factor prices in 1940.[22] This leaves us with 720 out of 744 CZs covering 99% of the U.S. residential population.

A.1.2 Variable Construction

Our main sample is households not living in group quarters whose highest earner is between 25 and 60 years old (i.e., the "representative household member"). For individual-level variables such as wage income or age, we use the values of the representative household member.[23]

To measure market wages (i.e., "full-time 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 "usual week" during the previous year (available 1980-2010) or during the previous week when "usual week" is not available (1940-1970). 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

[22] These two CZs contain Wayne county, UT with a population of 2,778 in 2010 and Norton county, KS with a population of 5,671 in 2010.

[23] 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).

wages.

To make homes comparable across decades, we use the sample selection applied by the Historical Census of Housing Tables.[24] We restrict homes (both owner-occupied and rental units) to non-farm single-family housing units on less than 10 acres without a business or medical office on the property.

Home values are self-reported and available for owners only. Condos were excluded in 1970 but included in 1980-2010. With the exception of 2010, the home value is not continuous and is reported as the midpoint of an interval. Except in 1940, estimates include the full value of both house and land, even if the respondent resides on only part of the property. In 1940, when data was also collected for single-family farmhouses, the estimate included only the value of the house, not the land. the value of owner-occupied units that were part of a building containing other households or businesses included only the part of the house in which the owner's household lived. For example, if the owning household of a two-family house rented half of the house to another household, only half of the house's value would have been reported. To adjust the home values in 1940, we scale up all values proportionately such that the national-level price-rent ratio matches the price-rent ratio 1960-1990 (excluding the boom bust period from the 1990s to the 2010s). Since this adjustment affects the level of all values in 1940 proportionately, it does not affect our measures of risk and hence any of our conclusions.

We use contract rent instead of gross rent because the latter is not available in the 1940 Census. Gross rent includes utilities whereas contract rent may or may not depending on the household. Figure C.1 shows that utilities are not or only weakly correlated with local income changes. Nevertheless, Figures C.6 and C.9 show that our results are robust to using gross rent and data from 1960 on.

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 log wage income on fixed effects for household year of age, and we regress log rents and log home prices on fixed effects for home age (5- and 10-year age bins), size (number of rooms) and other home characteristics, such at number of housing units (both occupied and vacant) in the structure containing the household, whether the home contains complete plumbing facilities, and whether it contains kitchen facilities. Housing characteristics are not available in 1940 and we therefore use the unadjusted values. Figures C.6 and C.9 show that we obtain similar results when dropping 1940 or

[24] https://www.census.gov/data/tables/time-series/dec/coh-values.html

when using unadjusted values over the entire sample period. All hedonic regressions also include location fixed effects and we add these fixed effects back to the regression residuals. Our CZ-level market price measures are thus stripped of differences in quantity and quality composition (e.g., age and size).

We winsorize variables at the 1st and 99th percentiles to account for changes in top and bottom coding thresholds an outliers and we use location-level medians instead of means, although we obtain similar results using means and other moments; see Figure C.3. All analyses use household sampling weights across households and population weights when looking across locations. All dollar values are converted to real 2020 dollars using the CPI - All Urban series.

B Wage, Rent, and Home Price Variation

This section provides extensive robustness checks of our main findings of the descriptive analysis of housing risk and provides details about the data generating process we estimate to compute the individual welfare results in the life-cycle model.[25]

B.1 Robustness of Relationships between Wages, Rents, and Home Prices

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 C.1 shows that location-level wage changes indeed have a negligible impact on these values.[26]

[25] Note that figures display log changes instead of growth rates.

[26] As a result, Panel (e) of Figure C.9 shows that the correlation of changes in wages and gross rents – which include utility costs – is attenuated relative to the correlation with market prices measured with contract rents.

B.1.2 Migration

Another concern is that these correlations might be the outcome of economically selected migration and hence a reflection of increased cross-location economic segregation. This might be the case for instance if factor price changes were largely driven by migrants and the correlations would therefore not reflect the net income risk of non-moving residents. However, Figure C.2 shows that the correlations are largely unchanged when we control for population changes or when we drop migrants from the sample.[27]

B.1.3 Housing Market Segmentation and Differences Across the Distribution

Related to the previous concern, we also show that these correlations apply across the entire cross-sectional distribution of wages and housing costs. Figure C.3 shows that the 25th percentiles and the 75th percentiles of the wage and housing cost distribution move similar to the median, and similarly the mean (after making top and bottom codes consistent over time by winsorizing variables at the 1^{st} and 99^{th} percentile).

Housing market for different types of homes might be segmented within a location. 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, Figure C.4 limits the sample of housing units to single-family homes (panels a and b), to homes of modal size of 6 rooms (panels c and d). The figure shows that the correlations between wages and housing costs are very similar as when we use the full sample of housing units.

Panels (e) and (f) extend the hedonic regressions to control for differential changes in the quantity and quality of the housing stock within location over time by projecting off changes in average home characteristics and average worker age within CZ over time. The resulting relationship between wages and housing costs is similar to the baseline results in Figure 2.

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

B.1.4 Property Taxes

Figure C.5 shows that that property tax rate changes – an important component of carrying costs – are uncorrelated with local income changes. We use this evidence and the fact that other costs such as utilities also do not correlate with local income (Figure C.1) to justify excluding carrying costs from our main analysis.

B.1.5 Robustness to Dropping 2000s Housing Boom and 1940 Census

Panels (a) and (b) of Figure C.6 drop the 2000s to test whether the wage and housing costs correlations are driven by the recent housing boom and bust cycle and the Great Recession. Excluding this decade does not affect the correlations much.

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. Panels (c) and (d) show similar results after dropping the 1940 Census and limiting the sample to 1960-2010, thereby dropping much of the post-war economic boom.

B.1.6 Alternative Definition of Local Labor and Housing Markets

Figure C.7 shows that the correlations are robust to using alternative geographic definitions of labor and housing markets. The strength of the correlation between wages and housing costs is similar at the county and the CZ level.

B.1.7 Self-Reported Survey Home Values versus Housing Transaction Data

Figure C.8 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.1.8 Alternative to Hedonic Regressions: Not Adjusting for Age and Size

Panels (a) and (b) of Figure C.9 show that not adjusting the for home characteristics and worker has little effect on the estimated correlations. While the hedonic regression substantially reduces the cross-sectional dispersion of full-time wages, home prices and rents, especially within locations, it as little effect on their *growth rates*.

B.1.9 Using Family Income instead of Wages

Panels (c) and (d) of Figure C.9 extends the 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 similarly correlated with housing costs as are full-time wages. 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 social insurance such as unemployment insurance, old-age pensions, or disability insurance).

B.1.10 Large and Small Commuting Zones

Figure C.10 shows that the relationship between wages and housing costs tends to be stronger in larger CZs.

B.2 Extensions of Empirical Analysis

B.2.1 Idiosyncratic Risk

[To be added]

B.2.2 National-Level Risk

[To be added]

B.2.3 Carrying Costs

[To be added]

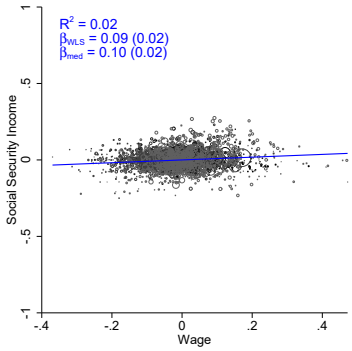
B.3 Data Generating Process (DGP) for Life Cycle Model

[To be added]

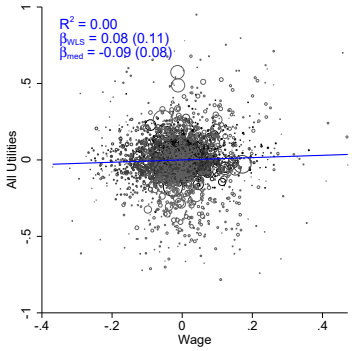
C Life Cycle Model

[To be added]

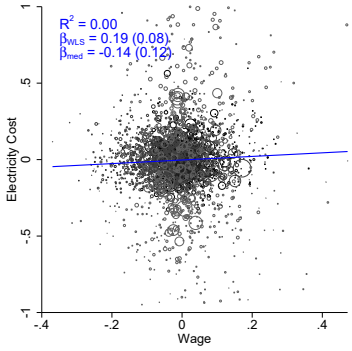
Figure C.1: (Non-)Correlation of Wages, Social Security, and Utility Costs



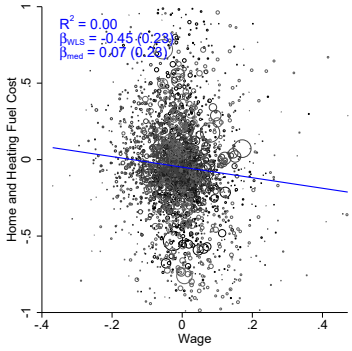
(a) Social Security/income



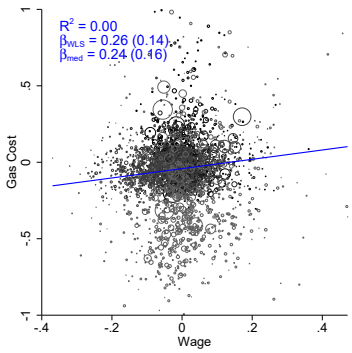
(b) all utilities/income



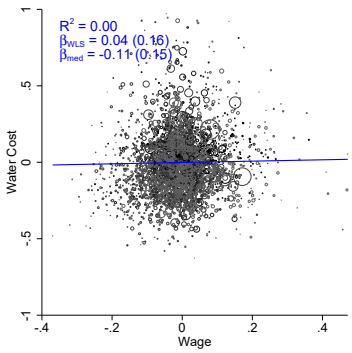
(c) electricity cost/income



(d) fuel cost/income



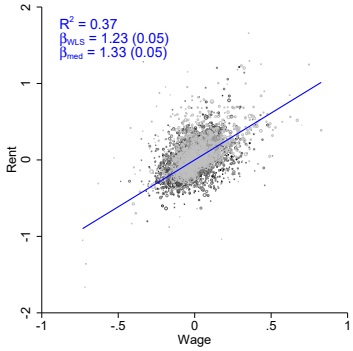
(e) gas cost/income



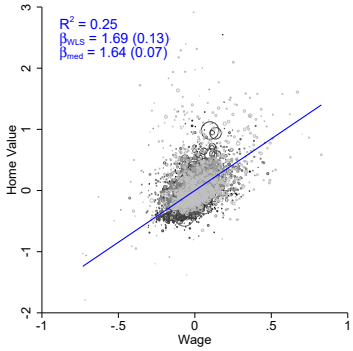
(f) water cost/income

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. Log changes of medians by Commuting Zone (CZ). (a) use Social Security income; (b) to (f) use utility expenses and its components, which are available from 1970 on. Circle areas represent CZ population size relative to total population. Observations in darker shades are from more recent decades. Source: Decennial Census and ACS.

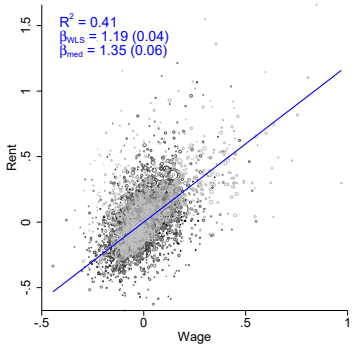
Figure C.2: Robustness to Controlling for Pop. Changes and Dropping Movers



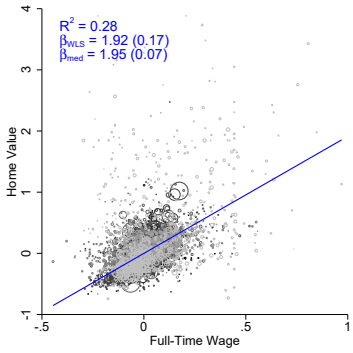
(a) rent/wage: population change



(b) price/wage: population change



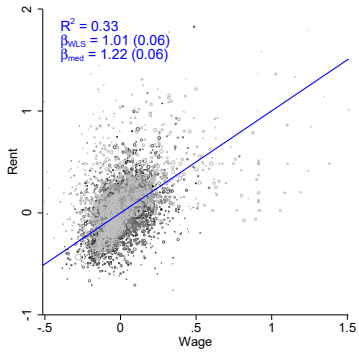
(c) rent/wage: dropping movers



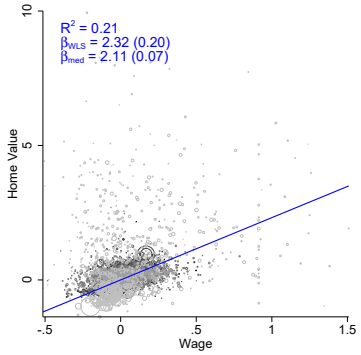
(d) price/wage: dropping movers

Notes. Log changes of median full-time wages, rents and home values by Commuting Zone (CZ). (a) and (b) project off log changes of CZ population (full Census counts) and CZ-level log changes in the number of full-time workers, rental units and owner-occupied homes in the sample. (c) and (d) drop respondents who report having moved between counties within the past 5 years (respectively 1 year for the 2010 sample). Circle areas represent CZ population size relative to total population. Observations in darker shades are from more recent decades. Source: Decennial Census and ACS.

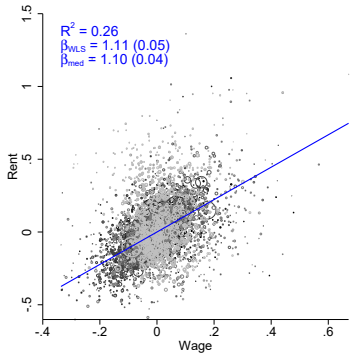
Figure C.3: Robustness across Distribution within Location



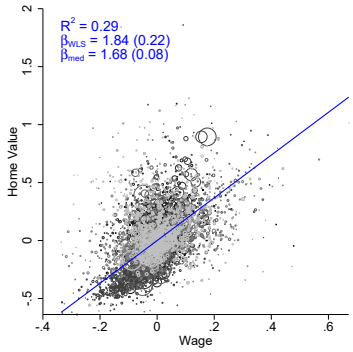
(a) rent/wage: 25th percentile



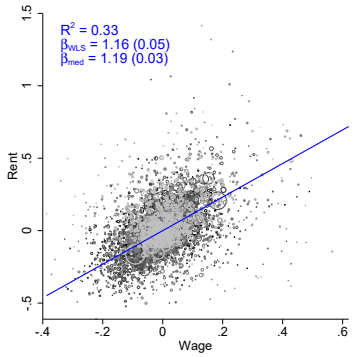
(b) price/wage: 25th percentile



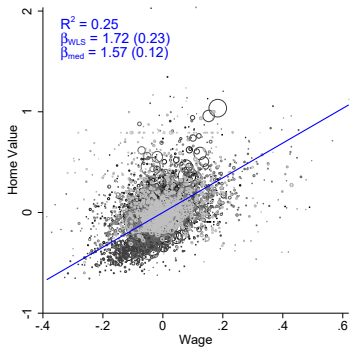
(c) rent/wage: mean



(d) price/wage: mean



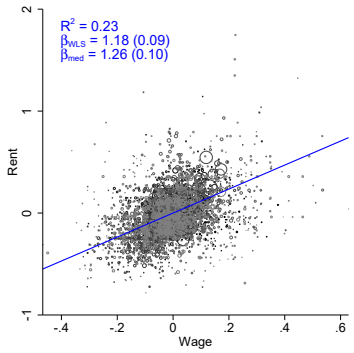
(e) rent/wage: 75th percentile



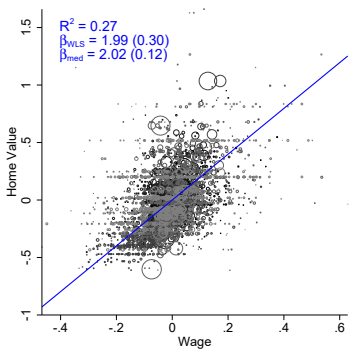
(f) price/wage: 75th percentile

Notes. Log changes of the mean (panels c and d), 25th (a and b) and 75th percentile (e and f) of full-time wages, rents and home values by Commuting Zone (CZ). Circle areas represent CZ population size relative to total population. Observations in darker shades are from more recent decades. Source: Decennial Census and ACS.

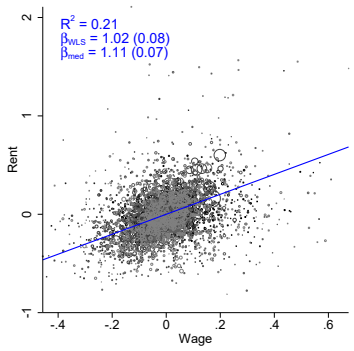
Figure C.4: Restrict to Single-Family Homes, 6-Room Homes, and Controlling for Composition Changes



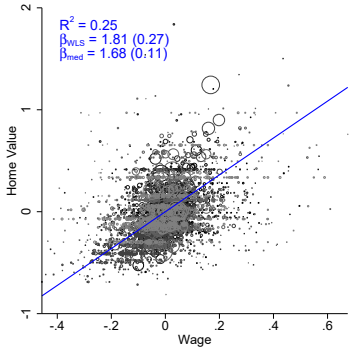
(a) rent/wage: single-family homes



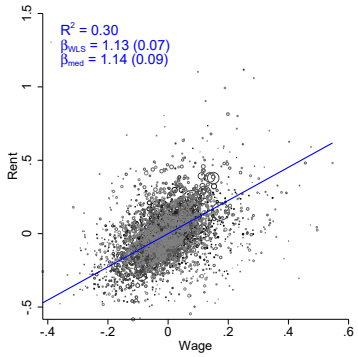
(b) price/wage: single-family homes



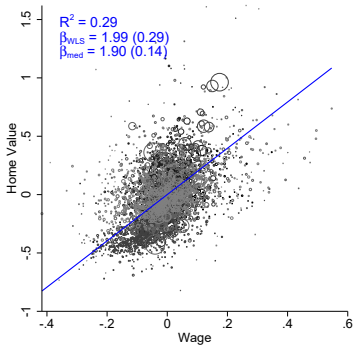
(c) rent/wage: modal home size



(d) price/wage: modal home size



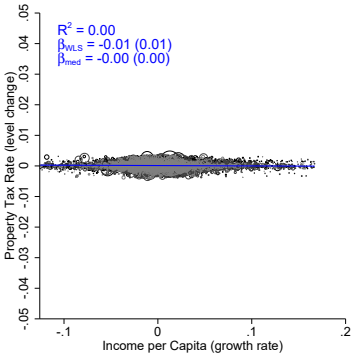
(e) rent/wage: control for characteristics



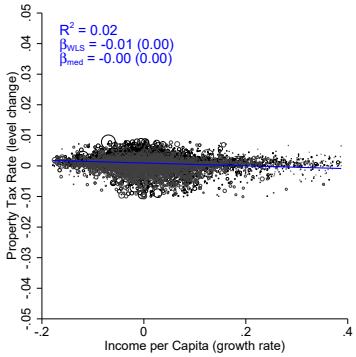
(f) price/wage: control for characteristics

Notes. Log changes of median full-time wages, rents and home values by Commuting Zone (CZ). (a) and (b) restrict the sample to detached single-family houses (the modal housing unit time) and use not hedonically adjusted values; (c) and (d) restrict the sample to 5-room homes (the modal number of rooms) and use not hedonically adjusted values; (e) and (f) control for composition changes by projecting off changes in average home characteristics and average worker age within CZ over time. Circle areas represent CZ population size relative to total population. Observations in darker shades are from more recent decades. Source: Decennial Census and ACS.

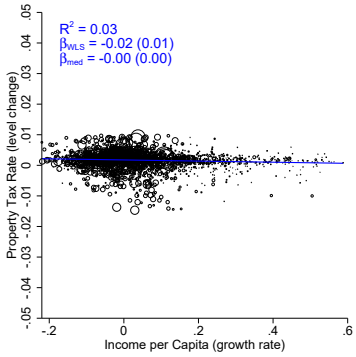
Figure C.5: (Non-)Correlation of Income and Property Tax Rates



(a) 1-year growth



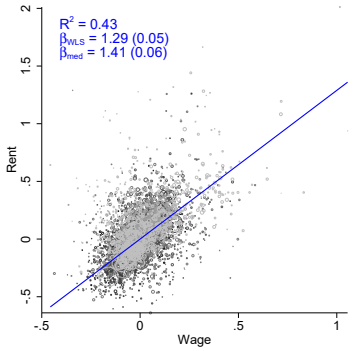
(b) 5-year growth



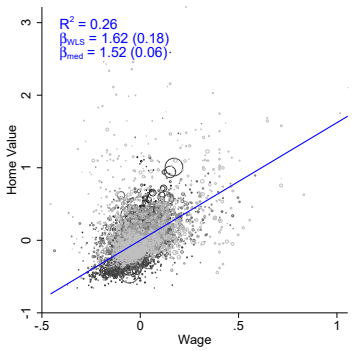
(c) 10-year growth

Notes. Log changes of per capita personal income and property tax rates by county, 2003-2015. All variables have been demeaned relative to national level log changes. (a) uses 1-year log changes; (b) uses 5-year log changes; and (c) uses 10-year log changes. Circle areas represent county population size. Observations in darker shades are from more recent years. Source: Baker et al. (2020) and Bureau of Economic Analysis.

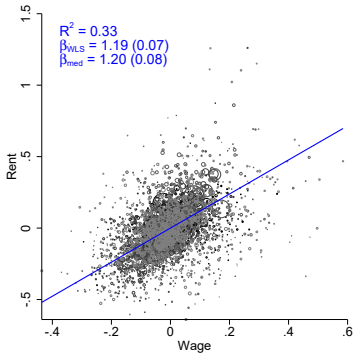
Figure C.6: Robustness to dropping 2000 Housing Boom and 1940 Census



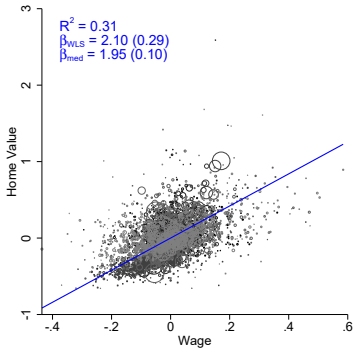
(a) rent/wage: drop 2000s housing boom



(b) price/wage: drop 2000s housing boom



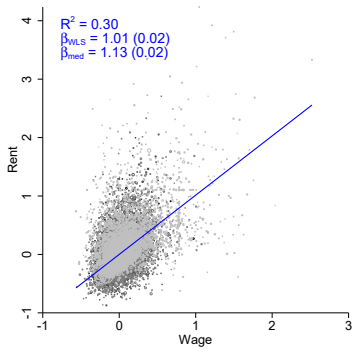
(c) rent/wage: drop 1940 Census



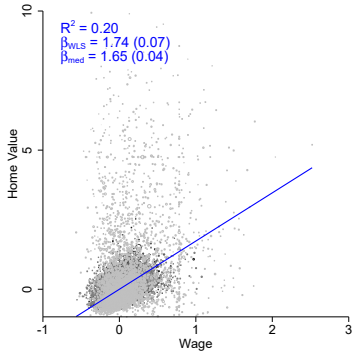
(d) price/wage: drop 1940 Census

Notes. Log changes of median full-time wages, rents and home values by Commuting Zone (CZ). (a) and (b) drop 2010 and hence the housing boom in the 2000s; (a) and (b) drop the 1940 Census. Circle areas represent CZ population size relative to total population. Observations in darker shades are from more recent decades. Source: Decennial Census and ACS.

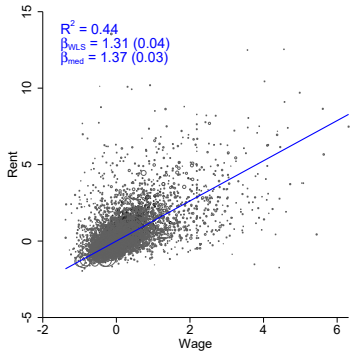
Figure C.7: Robustness to using Counties



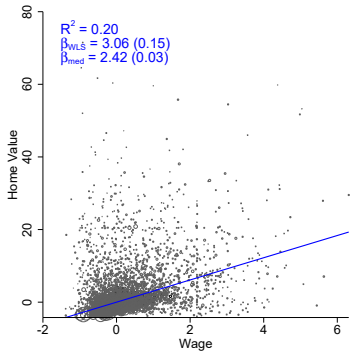
(a) rent/wage: 10-year growth



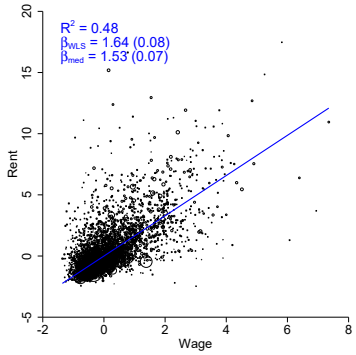
(b) price/wage: 10-year growth



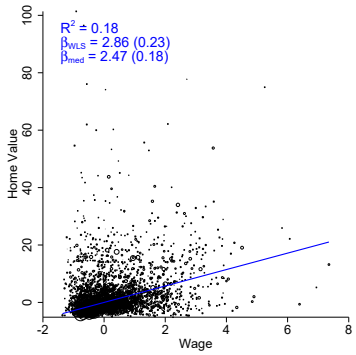
(c) rent/wage: 40-year growth



(d) price/wage: 40-year growth



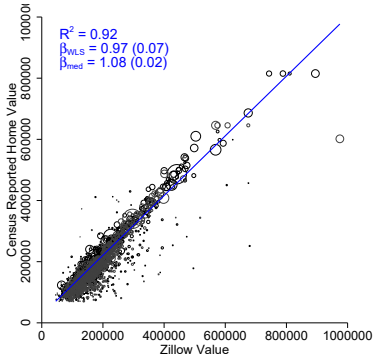
(e) rent/wage: 70-year growth



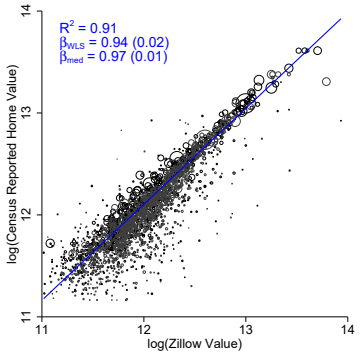
(f) price/wage: 70-year growth

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

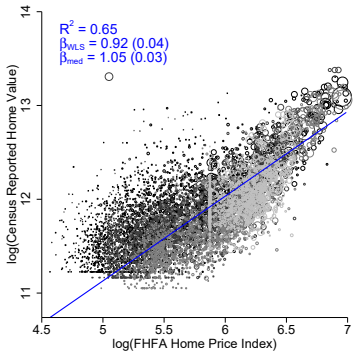
Figure C.8: Comparison with FHFA and Zillow Transaction Prices



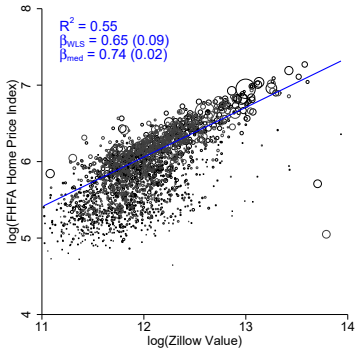
(a) Census/Zillow: levels



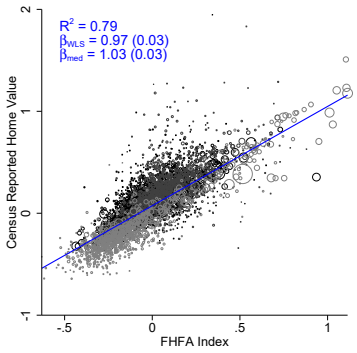
(b) Census/Zillow: logs



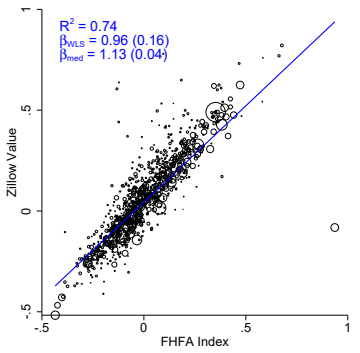
(c) Census/FHFA: logs



(d) FHFA/Zillow: logs



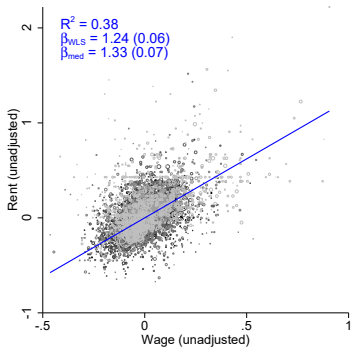
(e) Census/FHFA: 10-year log changes



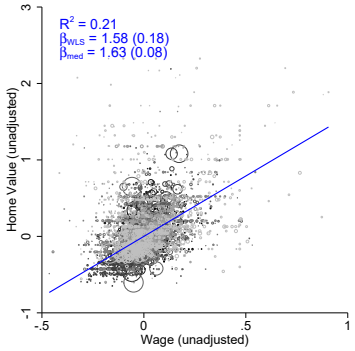
(f) FHFA/Zillow: 10-year log changes

Notes. Census home value (in dollars), Zillow Home Value Index (ZHVI) All Homes Time Series (in dollars), and Federal Housing Finance Agency (FHFA) Annual House Price Index (index value) by county. Circle areas represent county population size relative to total population. Observations in darker shades are from more recent decades. Source: Decennial Census, ACS, Zillow and FHFA.

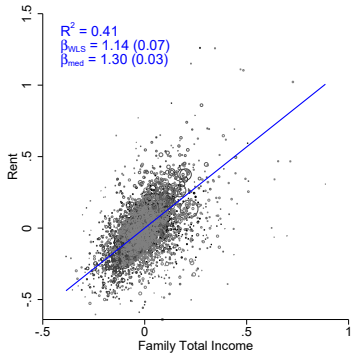
Figure C.9: Using Unadjusted Values, Family Income & Gross Rent



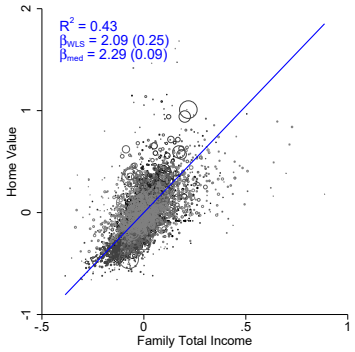
(a) rent/wage: unadjusted values



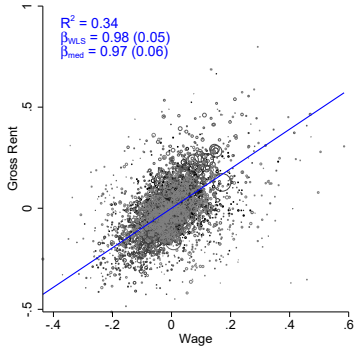
(b) price/wage: unadjusted values



(c) rent/family income



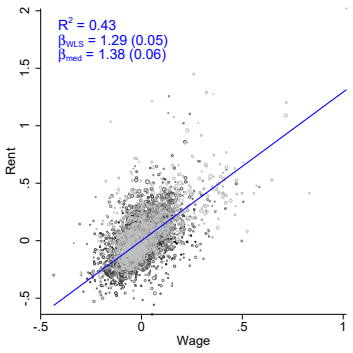
(d) price/family income



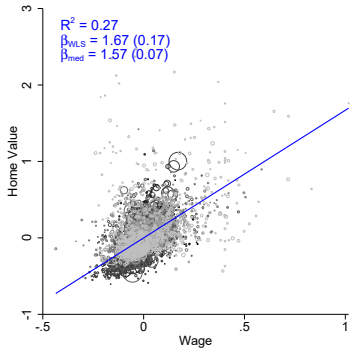
(e) gross rent/wage

Notes. Log changes of medians by Commuting Zone (CZ). (a) and (b) use raw values not hedonically adjusted for differences in home and worker characteristics. (c) and (d) use total family income instead of full-time wages. (e) uses gross rent, which includes utility expenses and is available from 1960 on, instead of contract rent. Circle areas represent CZ population size relative to total population. Observations in darker shades are from more recent decades. Source: Decennial Census and ACS.

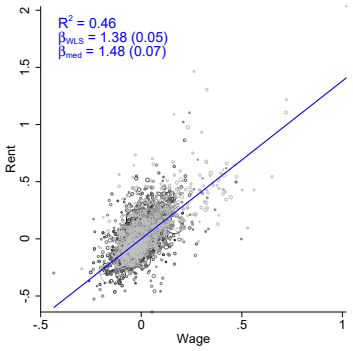
Figure C.10: Effect of Dropping Small Locations (CZs)



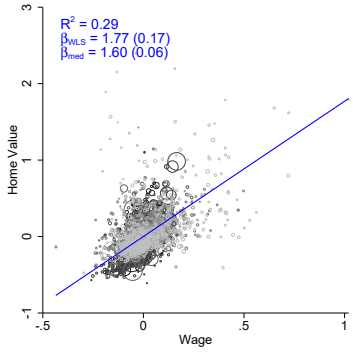
(a) rent/wage: CZ size > 25th percentile



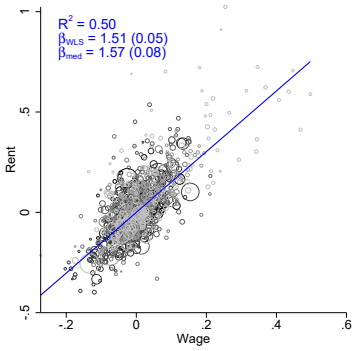
(b) price/wage: CZ size > 25th percentile



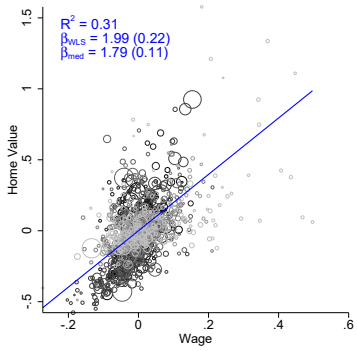
(c) rent/wage: CZ size > 50th percentile



(d) price/wage: CZ size > 50th percentile



(e) rent/wage: CZ size > 75th percentile



(f) price/wage: CZ size > 75th percentile

Notes. Log changes of median full-time wages, rents and home values by Commuting Zone (CZ). (a) and (b) drop CZs with a population below the 25th percentile, which corresponds approximately to a population size of 10,000 in 2010; (c) and (d) drop CZs with a population below the 50th percentile, which corresponds approximately to a population size of 50,000 in 2010; (e) and (f) drop CZs with a population below the 75th percentile, which corresponds approximately to a population size of 100,000 in 2010. Circle areas represent CZ population size relative to total population. Observations in darker shades are from more recent decades. Source: Decennial Census and ACS.