

Does Home Owning Smooth the Variability of Future Housing Consumption?*

Andrew Paciorek[†] and Todd Sinai[‡]

September 30, 2011

Abstract

We show that the hedging benefit of owning a home reduces the variability of housing consumption after a move. When a current home owner's house price covaries positively with housing costs in a future city, changes in the future cost of housing are offset by commensurate changes in wealth before the move. Using Census micro-data, we find that the cross-sectional variation in house values subsequent to a move is lower for home owners who moved between more highly covarying cities. Our preferred estimates imply that an increase in covariance of one standard deviation reduces the variance of subsequent housing consumption by about 11 percent. Households at the top end of the covariance distribution who are likely to have owned large homes before moving get the largest reductions, of up to 40 percent relative to households at the median.

1 Introduction

With the median U.S. family devoting about one-third of its annual income and 45 percent or more of its net worth to housing, fluctuations in house prices and annual housing costs have the potential to generate significant consumption volatility. Most analysts have focused on the effects on households of the sizable year-to-year fluctuations in house prices within a housing market. However, the effects of housing cost volatility may be mitigated merely by owning one's house (Sinai and Souleles 2005). Instead, a potentially significant source of housing cost uncertainty is faced by households who anticipate moving to different housing markets. While the average

*We thank Gadi Barlevy, Fernando Ferreira, Joe Gyourko, Chris Paciorek, Stuart Rosenthal, David Rothschild, Albert Saiz, Stephen Shore, Nicholas Souleles, Joel Waldfoegel, Justin Wolfers, Maisy Wong, two anonymous referees and seminar participants at Wharton, the Duke ERID housing dynamics conference, and the Federal Reserve Bank of Chicago for helpful comments and suggestions. We are grateful to the Research Sponsors Program of the Zell-Lurie Real Estate Center at Wharton for funding. The views we express in this paper are not necessarily those of the Federal Reserve Board of Governors or its staff. All errors are our own.

[†]Federal Reserve Board of Governors. E-mail: andrew.d.paciorek@frb.gov

[‡]Real Estate Department, Wharton School, University of Pennsylvania. E-mail: sinai@wharton.upenn.edu

standard deviation in real annual house price growth within a housing market is 5.6 percent, the differential growth in housing costs *across* markets has a standard deviation of 7.4 percent.¹ Thus a chance of moving to a new housing market creates uncertainty about the future price of housing, and the act of relocating could induce volatility in housing and non-housing consumption due to unanticipated differences in housing costs.

In this paper, we show that simply owning a house in the present can partially insure a household against uncertain housing costs due to potential moves in the future. The reason is that households tend to move between housing markets with correlated house prices, so their current houses often are worth more precisely when their next house is more expensive. This positive correlation between wealth and house prices mitigates the decline in housing consumption due to higher prices alone, and dampens the increase in housing consumption due to lower prices.

We illustrate this idea in a simple two-period representative agent model with migration and stochastic house prices. The model predicts that housing consumption for households who move between housing markets with higher covariances should vary less in the cross-section, since unexpected price shocks in the destination market are more likely to be matched by price changes in the origin. For example, households who move between highly covarying cities, such as New York and Boston, typically experience offsetting wealth and price effects on housing demand, as they are likely to face high house purchase prices in their destination cities when their sale prices in their origin cities are high, and vice versa. Conversely, households who move between low covariance city pairs, such as New York and Dallas, are relatively more likely to be wealthier when house prices in their destinations are *lowest*, and vice versa.

If the income and price elasticities of demand for housing have opposite signs, as is commonly believed, there should be less variance in subsequent housing consumption across households who move between city pairs like New York/Boston than among households who move between markets similar to New York/Dallas, holding the distributions of house price shocks constant. The model also predicts that the hedging benefit of high covariance would be strongest for households who

¹The figure of 7.4 percent is the standard deviation in the difference between the annual house price growth in one's own housing market versus other markets. Calculated using the Federal Housing Finance Agency's conventional mortgage repeat sales index, deflated by the CPI, for 168 metropolitan areas over the 1982-2007 time period.

own more housing before moving, since these households have more invested in the hedging asset.

Using a cross-section of household-level microdata from the U.S. Census, we find empirical support for these predictions. In particular, we find a negative relationship between the conditional variance in post-move house values across home owners who recently moved and the covariance in the house prices of their origin and destination cities. Our two-stage conditional variance procedure includes controls at both stages of the estimation for household demographics and the expected hedging benefit of home ownership, as well as origin and destination Metropolitan Statistical Area (MSA) fixed effects.² The demographic controls account for predictable differences in the level of housing demand, variation in the demographic composition of movers across MSA pairs, and any heteroskedasticity related to observable household characteristics. The MSA fixed effects account for differences among MSAs that could affect the level of housing wealth or spending, such as the level of house prices, or the variance, such as households departing an origin MSA having systemically higher variance in wealth.

Because of the origin and destination MSA fixed effects, our estimates of the hedging benefit are identified primarily from the pairing of MSAs. For example, in any given year, each household who moves out of New York receives a common house sale price shock, with those who move to Boston receiving a different house purchase price realization than those who move to Dallas. Also, households moving out of San Francisco in that year receive a *different* house sale price shock, with some of those households moving to Boston (which has high covariance with San Francisco) and others moving to Dallas (which does not). By pooling all such MSA origin/destination combinations, we can identify how the conditional variance of housing demand varies with city-pair house price covariance while still holding constant the characteristics of the origin and destination cities. In addition, households in our data move at any time during a five-year period, increasing the effective number of realizations of house price shocks for any origin/destination city pair.

Overall, we find that a one standard deviation increase in covariance reduces the variance of subsequent housing consumption for the average household by about 10 percent in our preferred

²Our two-stage estimation process is similar to that proposed by Engle (1982) to test for Autoregressive Conditional Heteroskedasticity (ARCH) disturbances in time series applications. It is also similar to Breusch and Pagan's (1978) test for heteroskedasticity.

specification. The reduction in variance is especially pronounced — as much as 18 percent — for home owners whom we predict were likely to own larger houses before moving. The total effect of covariance can be sizable for those households who move between highly covarying cities. An average household experiences nearly a 30 percent reduction in its variance of housing spending if it moves between the 95th percentile covarying city pair relative to a move at the median. For a high-income family, the same comparison yields a reduction in variance of 40 percent.

We find two other pieces of evidence consistent with the hedging interpretation of the relationship between covariance and the subsequent variance in housing spending. First, the theoretical prediction that the housing hedge should have the largest effect for households who owned more housing before moving enables us to use household-level variation to relax the identifying assumption of the same (conditional) distribution of housing demand for all movers from an origin city or to a destination city, and instead allow that distribution to vary by MSA pairs. We show that, among households who move between a origin-destination city pair, the effect of covariance on the conditional variance in subsequent housing spending is largest for households who had high incomes relative to prices in the origin or had high predicted housing consumption in the origin MSAs. Second, households with higher covariances are more likely to own a house in their destination MSAs, consistent with the hedge ensuring that they have enough wealth to buy a home.

Finally, we use a generalized additive model (GAM) to allow for a nonlinear effect of covariance in house prices on the variance (and mean) of housing consumption after a move. The GAM also allows us to non-linearly control for covariates such as household income and age. The estimated relationship between variance and covariance is noticeably nonlinear, with a larger effect on the margin for households with high covariance. For example, the relationship for households below the mean of covariance is about one third as large as for those above. This indicates that high-covariance households have a better hedge on the margin, as well as on average.

This paper makes contributions in several areas. First, previous research has estimated how the hedging potential of home ownership affects households' *ex ante* choices of tenure mode or the quantity of housing to consume (Sinai and Souleles 2005, Han 2008, Sinai and Souleles 2009), or has considered the hedging properties of home ownership in theory (Ortalo-Magne and Rady 2002).

Other researchers have noted that, due to its nature as a consumption commitment and illiquid asset, owning a home shifts households' financial investment choices (Cocco 2004) and also affects the volatility of consumption by changing how consumption responds to income shocks (Chetty and Szeidl 2007). In contrast, this paper shows that owning a home does reduce the *ex post* variance of housing consumption, so that households are correct in believing that owning a house can hedge future housing costs.

Second, this paper contributes to the consumption smoothing literature by providing an important example of what Cochrane (1991) calls an “informal institution” that provides consumption insurance. Consumption smoothing has been examined in a number of contexts, such as unemployment insurance (Gruber 1997, Browning and Crossley 2001, Chetty and Szeidl 2007) and welfare (Gruber 2000), typically in the sense that there are institutions that facilitate a consistent level of consumption when there are unexpected changes in income. In our context, owning a home enables a household to better maintain a level of housing consumption in the face of unexpected changes in prices. Hurst and Stafford (2004) and Hryshko, Luengo-Prado and Sorensen (2010) consider the effect of the liquidity provided by housing equity on the smoothing of *non*-housing consumption. While we do not provide direct empirical evidence on non-housing consumption, if owning a home hedges future housing consumption it should also reduce the variance of non-housing consumption, since the entire consumption bundle is affected by changes in house prices.

Our empirical approach is in the spirit of Cochrane (1991), who regresses a household's change in consumption on a proxy for an idiosyncratic shock to income, such as illness, and Gruber (1997), who compares how consumption responds to unemployment when unemployment insurance is more or less generous. There are two important distinctions between our context and these papers: First, our shocks are to relative prices (of housing) rather than income. Second, we implement a more complex empirical strategy that infers the consumption response to a shock from the cross-sectional conditional variance in housing consumption subsequent to a move. We take this approach not only because we do not observe in our data housing consumption prior to moving, but even if we did, the durable nature of housing implies that the amount of housing consumption prior to moving would be a poor proxy for the the latent desired housing consumption after moving.

Our discussion proceeds as follows: In the next section, we outline a simple model of housing consumption with migration. We use this model to derive the direct effect of the covariance of house price changes on the variance of subsequent housing purchases, and discuss how the response of the initial housing choice to that covariance can induce a second-order indirect effect. In Section 3, we discuss our data and detail our use of cross-sectional and time-series variation to identify the effect of the covariance hedge. Then, in Section 4, we describe our conditional variance empirical strategy. We present our results in Section 5 and interpret the magnitudes in Section 6. Section 7 concludes.

2 A Simple Model of Housing Consumption with Migration

2.1 Intuition

In this section, we focus on what happens when, because it previously owned a home, a household's wealth is not independent of the house prices it faces after a move. Since a large fraction of household wealth is allocated to housing, if house prices in the market a home-owning household is moving *from* covary positively with house prices in the market the household is moving *to*, the household will be wealthier (due to selling its prior house) precisely when the next house is more expensive and poorer when the next house is relatively cheap. Standard housing demand models recognize that households with more wealth should buy more housing, all else equal, and those that face higher prices should buy less. These wealth and price effects on housing demand will offset each other for those households for whom house prices in the former and next housing markets covary more strongly, and especially for those who have allocated more of their wealth to housing, thus providing a natural hedge against house price volatility.

This hedging intuition suggests that the potential variance in housing consumption in the destination market should be lower among households who moved between more highly covarying housing markets than for those who moved between more independent housing markets. For high covariance households, the effect on housing demand of the varying housing prices they face — due to moving between different markets or from moving at different points in time — would be undone

more by the effect on housing demand of their housing wealth.

An example of this intuition is demonstrated in Table 1, where we consider the effect on housing demand of the polar cases of perfectly positive (negative) covariance in house price growth between the origin city, A, and destination city, B. For the sake of the example, we assume that a household's wealth is entirely made up of their house in city A, and we use as parameters three sets of estimates of the elasticity of housing demand with respect to wealth (ϵ_w) and house prices (ϵ_p).

In the first panel of Table 1, we consider the case of perfect covariance — something like the New York-Boston example described above. In that case, if the household faces 20 percent growth in house prices in city B, which would reduce its demand for housing there, house prices also grow commensurately in city A, making the household 20 percent wealthier and raising its demand for housing in city B. The first row of column (4) shows that under Cobb-Douglas preferences with wealth and own-price demand elasticities of 1 and -1 respectively the price and wealth effects would net out exactly and there would be no change in housing consumption after the move. The second row shows that if house prices in city B *decline* by 20 percent, there still is no effect on post-move consumption. With perfect covariance, the 20 percent decline in house prices in B implies a 20 percent decline in wealth and under Cobb-Douglas preferences the price and wealth effects exactly offset.

Under other plausible elasticity estimates, the wealth and price elasticities are not equal and so housing consumption in city B responds to the price change even in the perfect covariance case. The last two columns illustrate this point using two widely differing pairs of elasticities from Rosen (1979) and Borsch-Supan (1990).³ Using Rosen's (1979) estimates, consumption after the move is five percent lower in the simultaneous increase case shown in the first line, and five percent higher in the decrease case shown in the second. Under Borsch-Supan's (1990) elasticities, the figures are negative and positive four percent, respectively.

By contrast, the bottom two rows of Table 1 considers the case where house price growth is

³Both papers estimate demand elasticities with respect to income rather than wealth; we use them here simply as examples. Estimates of housing demand elasticities vary widely in the literature. See, for example, Glaeser, Kahn and Rappaport (2008), Haurin, Hendershott and Kim (1994) and Hansen, Formby and Smith (1998). We require only that housing demand be strictly increasing in wealth or permanent income and strictly decreasing in price, all else equal.

perfectly negatively covarying in cities A and B, more akin to our earlier New York-Dallas example. Thus, a 20 percent rise in house prices in city B is accompanied by a 20 percent decline in house prices in city A, and vice versa. In this case, the price effect and wealth effect work in the same direction. Under Cobb-Douglas parameters, housing demand in city B falls by 40 percent when prices rise in B or grows by 40 percent when prices fall in B. Under the parameters estimated in Rosen (1979), the effect on housing demand is -35 percent or $+35$ percent, respectively. Even when the elasticities are much smaller, as in Borsch-Supan (1990), the effect is negative 8 percent in the third line and positive 8 percent in the fourth.

The intuition of the paper can be seen by comparing the positive covariance and negative covariance cases. Regardless of which elasticities we use, the potential variation in housing demand in city B is greater for the negative covariance households. This result follows from the wealth and price elasticities having opposite signs, so when the covariance is positive, the household's exposure to volatility in prices in city B is hedged by wealth changes due to co-movements in price in city A. When the covariance is negative, the household suffers from a negative hedge, so that their wealth is lowest precisely when prices in B are highest.

2.2 Model setup

To generalize this intuition and provide guidance for the econometric specification, we outline a simple dynamic consumption model. Since our focus is on the relationship between house prices in current and future housing markets, our model encompasses multiple locations with stochastic house prices but abstracts away from other complications.⁴ In particular, we ignore transactions costs from moving and keep the decision of when and where to move exogenous and known ex ante; we discuss the implications of endogenous moving in the section on empirics below.

We consider the decisions of an infinitely-lived agent. At the beginning of time period t , he moves to city m_t from city m_{t-1} . He enters the period with wealth w_t and receives (exogenous and

⁴For example, Davidoff (2006) allows for the covariance of labor income and housing costs; Shore and Sinai (2010) consider the effect of the fixed cost of moving; Sinai and Souleles (2005) endogenize house prices in the origin city; Piazzesi, Schneider and Tuzel (2007) allow for correlations between the returns of housing and other assets; Ortalo-Magne and Prat (2009) endogenize house prices and portfolio returns in a multi-city model; and Davidoff (2010) allows for changes in the marginal utility of housing to be correlated with the marginal utility of long-term care.

known ex ante) labor income l_t , the sum of which he must divide between non-housing consumption (c_t), housing investment (h_t) and investment in financial assets (s_t). There is no rental sector in the model, so the agent must purchase a home in m_t that costs $h_t = P_t^{m_t} q_t$ for a house of size q_t . For simplicity, we let one unit of housing produce one unit of housing services and define the utility function accordingly. We denote per-period utility over numeraire consumption and housing services by $u(c_t, q_t)$.

At the end of period t , the agent receives stochastic returns of $1 + r_{t+1}$ per dollar invested in financial assets and $1 + \pi_{t+1}^{m_t} = \frac{P_{t+1}^{m_t}}{P_t^{m_t}}$ per dollar invested in housing.⁵ These sum to his next-period wealth w_{t+1} , so that

$$w_{t+1} = q_t P_t^{m_t} (1 + \pi_{t+1}^{m_t}) + s_t (1 + r_{t+1}) \quad (1)$$

The dividend portion of the housing return, the rental value, is consumed in-kind by living in the house. House price growth may be correlated — positively or negatively — across markets, but we assume that house price growth rates in all markets are uncorrelated with financial returns.⁶ We also assume that it is not possible to go long or short either housing market except through the purchase of a home, so investment in housing cannot be divorced from the consumption of housing services.

House prices and financial returns are stochastic, so the agent must form (rational) expectations about the future and maximize over consumption, housing, and financial investment accordingly. He discounts next period's utility by a factor of β . Since the agent has an infinite horizon and stable preferences, we can easily write his problem using the Bellman equation. He solves

$$V(w_t; P_t, S_t) = \max_{c_t, q_t, s_t} \{u(c_t, q_t) + \beta E[V(w_{t+1}; P_{t+1}, S_{t+1}) | P_t, S_t]\} \quad (2)$$

⁵We do not restrict s_t to be positive, which allows the agent to borrow to finance housing consumption. For example, if s_t is negative and r_{t+1} is known with certainty, the “financial asset” acts like a fixed-rate mortgage. Since households must borrow or lend at the same rate of return and taxes are ignored in the model, there is no reason to borrow to finance the purchase of housing while simultaneously investing in a financial asset. Extending the model to allow for this common behavior does not materially affect the analysis.

⁶Flavin and Yamashita (2002) show that house price growth in the Panel Study of Income Dynamics (PSID) had a correlation with the S&P 500 of nearly zero over the period 1968-1992. We calculate that house price growth rates in more than 90 percent of metropolitan statistical areas have correlations with stock returns of between -0.2 and 0.2.

subject to the budget constraint

$$c_t + q_t P_t^{m_t} + s_t = w_t + l_t$$

The state space S_t can contain any information available at time t that is useful for predicting future returns. We write P_t , the full vector of house prices in all cities at time t , separately to emphasize the importance of house prices in our model.

We make a set of standard assumptions about preferences in order to analyze the model. First, we assume that the per-period utility function $u(c_t, q_t)$ is twice continuously differentiable, strictly increasing and strictly concave in both consumption and housing. Consequently, the value function is twice continuously differentiable, strictly increasing and strictly concave, so that the agent is risk-averse with respect to post-move wealth. Second, we assume that $u_c(c_t, q_t)$, the derivative of the utility function with respect to non-housing consumption, is twice differentiable and itself has a positive second derivative with respect to consumption. This implies that the agent would be a precautionary saver if his housing purchase were fixed.⁷

Under these assumptions, it is straightforward to derive Euler equations that define the agent's choices of q_t , s_t and c_t . They are

$$-u_c(c_t, q_t) P_t^{m_t} + u_q(c_t, q_t) + \beta E [u_c(c_{t+1}, q_{t+1}) P_t^{m_t} (1 + \pi_{t+1}^{m_t})] = 0 \quad (3)$$

$$-u_c(c_t, q_t) + \beta E [u_c(c_{t+1}, q_{t+1}) (1 + r_{t+1})] = 0 \quad (4)$$

where letter subscripts denote the derivative of the function with respect to that argument.

2.3 Variance in the Marshallian Demand for Housing

In this subsection we derive theoretical predictions for the relationship between the covariance of pre- and post-move house prices and the variance in post-move housing consumption, holding pre-

⁷That is, if he were committed to his current house and could not adjust on that margin, a positive second derivative of the marginal utility of consumption would be necessary and sufficient to guarantee that post-move consumption would increase in response to an increase in the variance of post-move wealth. If housing consumption is not fixed, the precautionary saving motive can even be reversed (Shore and Sinai 2010). See Kimball (1990) and citations therein for a full discussion of the mathematics of the precautionary saving motive in the one-good case.

move housing investment, q_{t-1} , and all future decisions fixed.⁸ This relationship depends crucially on how q_t responds to changes in house prices in the origin ($P_t^{m_{t-1}}$) and destination ($P_t^{m_t}$). For housing to serve as a hedge, housing demand must be decreasing in destination house price and increasing in wealth (and thus origin house price).

It is possible to use equations 3 and 4 to derive exact conditions under which these statements are true in our model. We defer the math to Appendix A because the intuition is clear: Housing demand falls with higher prices, holding future prices constant, as long as housing is not a Giffen good. Demand rises in wealth as long as housing and consumption are complementary — or at least sufficiently non-substitutable — and the risk-adjusted return on the financial asset exceeds the risk-adjusted pecuniary portion of the return on housing in expectation, so that people are not tempted to over-consume housing just to save. This property of the returns follows from a no-arbitrage condition and the observation that part of the return on housing, the rental value, is consumed in-kind.

We proceed by examining the Marshallian demand function for housing in period t ,

$$q(w_t, P_t^{m_t}) = \arg \max_{q_t} \{u(c_t, q_t) + \beta E[V(w_{t+1}; P_{t+1}, S_{t+1}) | P_t, S_t]\} \quad (5)$$

which is again subject to the budget constraint. Taking a first-order Taylor approximation to $q(w_t, P_t^{m_t})$ at any point (\bar{w}, \bar{P}) yields

$$\begin{aligned} & Var [q(w_t, P_t^{m_t})] \\ & \approx Var [q(\bar{w}, \bar{P}) + q_w(\bar{w}, \bar{P})(w_t - \bar{w}) + q_p(\bar{w}, \bar{P})(P_t^{m_t} - \bar{P})] \\ & = (q_w(\bar{w}, \bar{P}))^2 Var [w_t] + (q_p(\bar{w}, \bar{P}))^2 Var [P_t^{m_t}] \\ & \quad + 2(q_w(\bar{w}, \bar{P}))(q_p(\bar{w}, \bar{P})) q_{t-1} Cov [P_t^{m_{t-1}}, P_t^{m_t}] \end{aligned} \quad (6)$$

The second equality simply applies the definition of w_t in Equation 1 and our assumption that financial returns are uncorrelated with housing returns.

The last two lines of equation 6 shows that post-move housing demand has higher variance

⁸We consider the implications of allowing q_{t-1} to vary in response to changes in variance or covariance in Section 2.4.

when variance in wealth is greater (the first term) or when the variance in destination house prices is greater (the second term). Since wealth comprises investments in financial assets and the origin house, greater variance in returns for either asset yields higher variance in the destination housing demand. This occurs because housing demand increases in wealth ($q_w(\cdot) > 0$) and decreases in price ($q_p(\cdot) < 0$), so volatility in either wealth or purchase price carries over into variance in housing demand.

Importantly, a higher cross-market price covariance reduces the variance in housing demand, all else equal. Any decrease in housing demand due to higher house prices at the destination are at least partially offset by the greater wealth from the higher sale price on the origin house, and vice versa.⁹ Equation 6 has immediate empirical implications. Conditional on origin and destination variance in house prices, higher covariance should yield a lower variance of housing demand. In addition, the q_t term multiplying the covariance term shows that the reduction in variance should be more pronounced for households who owned more housing in the prior period.

It is worth noting that greater covariance should also reduce the variance of *non*-housing consumption. Thanks to the income effect, when house prices rise, households will have less wealth to spend on non-housing consumption, and vice versa when house prices fall. Using the same intuition that underlies Equation 6, this income effect is offset when covariance is higher because origin house prices rise in tandem with destination house prices. This implication contrasts with Chetty and Szeidl (2007), who find that home ownership increases the sensitivity of consumption to income shocks. Our results apply to the variance in consumption around a move; Chetty and Szeidl's (2007) result holds when the commitment nature of home ownership precludes moving. While the smoothing of non-housing consumption is another important channel by which the covariance hedge can improve welfare, our data do not include information on non-housing consumption, so we will not be able to test it empirically.

⁹In this paper, we focus on whether owning a house hedges households along the dimension of having to substitute between housing and non-housing consumption. To that end, we treat housing as a composite good where the price of one unit of housing service flow can differ across MSAs or within MSAs over time but does not differ across types of housing within an MSA. In theory, it is possible that households also substitute among the characteristics that make up the housing bundle, such as land or structure. We have implicitly assumed either that the relative prices of the characteristics that make up the housing bundle do not change over time, or that the various dimensions must be consumed in constant proportions, or that households do not have preferences over the various characteristics.

2.4 Endogenizing Initial Housing Consumption

The analytical results above hold when wealth entering period t does not change in response to differences in covariance. However, housing consumption might respond to the expected benefit of home owning as a hedge. Supporting evidence has been found in prior research: Sinai and Souleles (2005) and Sinai and Souleles (2009) find that households are more likely to purchase a house if it is expected to provide a larger hedge, while Han (2008) and Han (2010) find theoretical and empirical evidence that households purchase more housing in that circumstance.

Our theoretical framework also generates this same marginal effect of covariance on the intensive margin of housing consumption in the period before the move ($t - 1$). One mechanism is that pre-move housing is more valuable when the expected hedge is stronger because it is better at reducing post-move consumption volatility. It is straightforward to show the basis for this intuition using our model of housing investment; the math is provided in Appendix A.

Another possible mechanism for an endogenous response of pre-move housing choices to covariance is due to the precautionary saving motive (Kimball 1990). Households whose houses provide a better hedge and thus are insured against future house price risk may choose to save less, spending more on both housing and non-housing consumption before moving. This suggests that the household will decrease financial saving s_{t-1} in favor of c_{t-1} and possibly q_{t-1} .

Regardless of the mechanism, when housing consumption and saving before the move respond to the hedging benefit of home ownership, there is a second channel — other than the direct hedging effect in Equation 6 — by which house price covariance can affect the variance of housing consumption in period t . Namely, households who enter period t with more wealth will experience more variance in consumption, all else equal, since they have more dollars in risky investments. Although this mechanism is not the focus of this paper, we will need to account for it in our empirical analysis. Consequently, we explicitly control for the effect of the expected covariance on w_t . We further explain this approach in Section 5.

3 Data

For our empirical work, we use the 5 percent sample from the 2000 Census Individual Public-Use Microsample (IPUMS) that contains household-level responses to the 2000 U.S. Census long-form questionnaire.¹⁰ We chose this data source because it reports a household’s MSA of residence in 2000 and in 1995, as well as household characteristics and housing spending in 2000.¹¹ From the two observations on the MSA of residence, we can infer moving between MSAs and subsequently match those moves to covariance in house prices across MSA pairs. Another benefit of the IPUMS data is that it contains enough observations that we can control nonparametrically for unobservable differences among origin or destination MSAs.

For house price data, we take the Federal Housing Finance Agency (FHFA) MSA-level house price indices, deflate by the Consumer Price Index, and peg them to the average house prices reported for each MSA in the 2000 Census. We use this real dollar-valued house price measure to calculate a cross-MSA covariance matrix.¹² These indices use repeat sales of houses with conventional mortgages to estimate constant-quality house price indexes for nearly 400 MSAs. Although the FHFA indices begin as early as the mid-1970s for some MSAs, we use the period starting in 1982 since more MSAs are available in the data starting at that time. We end the series in 1999 in order to use data prior to the observations in the 2000 IPUMS.¹³

In the IPUMS, we use all single-family, one- or zero-couple households that own a home, of which there are approximately 3.3 million in the sample.¹⁴ Of this group, we keep those who have moved from one MSA to another in the last five years and have both their current and previous MSAs of residence identified.¹⁵ This leaves a sample of about 150,000 households, across 284 origin

¹⁰Ruggles, Sobek, Alexander, Fitch, Goeken, Hall, King and Ronnander (2004)

¹¹An MSA is designed to correspond to a labor market area, and typically contains one or more focal cities and their surrounding suburbs.

¹²The results are similar, if less well grounded in our theoretical results, if we use covariances of real house price growth rates.

¹³The results are robust to changes in the horizon of the covariance calculation, such as using the whole period of data availability, from 1982 to 2007, or the period from 1995 to 2000 in which our households actually moved. Moreover, other house price series, such as those from S&P Case-Shiller or Core Logic, show very similar patterns of covariance over these time horizons to those in the FHFA data.

¹⁴This group excludes the few home-owning households with a household head — the first person listed on the Census form — under the age of 25.

¹⁵Since migration is reported for individuals and we do not want to explicitly account for household formation, we assume that the origin MSA of the household head is the origin MSA of the household. Unfortunately, the IPUMS

MSAs, 297 destination MSAs, and about 26,000 origin-destination pairs. Since the 17 years of FHFA data are not available for all MSAs identified in the IPUMS, this further limits our sample to about 100,000 households, 156 origin MSAs and 167 destination MSAs. Finally, to mitigate the effects of data reporting errors for house values or transitorily low incomes, we drop the top and bottom 1 percent of the observations based on their self-reported house price to income ratios and exclude any household with a MSA median house price-income ratio of above 10.

The summary statistics provided in Table 2 show how the observable characteristics change as the sample size diminishes. Average house value, household income, fraction married, and share college-educated rise, and average age declines. Most of the changes come from restricting the sample to movers and residents of MSAs, who tend to be better-educated and higher-income than rural residents. In the final sample, average house value is about \$230,000, household income is nearly \$100,000, and more than three-quarters of household heads are married.¹⁶

Figure 1 shows the distribution of house price correlations in our sample, weighted by the number of movers between each city pair. We have highlighted where ten representative MSA pairs lie on the correlation distribution. For example, Boston and New York have very highly correlated house prices (0.96), as do Dallas and Houston (0.85), and Chicago and Seattle (0.86). However, Dallas and New York have negatively correlated house prices (-0.26) and house prices in Chicago and Dallas are almost perfectly negatively correlated (-0.95). Sinai and Souleles (2009) show that, within any given origin MSA, there is wide variation across movers in the correlations of house prices in their destinations.

Our model implies that it is covariance rather than correlation that matters for ex-post housing variance. Intuitively, this is because when house prices have little volatility, there is less variance to hedge. Since it is possible for prices in two cities to be highly correlated but covary little if house price variance is low, we distinguish between correlation and covariance. Figure 2, which

does not identify all residents of many metropolitan areas. The IPUMS places the highest priority on identifying Public Use Microdata Areas (PUMAs). When PUMAs cross MSA boundaries, the MSA identifier may be suppressed to maintain maximum confidentiality. Still, most residents of most metropolitan areas are identified as such. See <http://usa.ipums.org/usa/volii/incompmetareas.shtml> for complete details.

¹⁶With the exception of household income, the mean values of these covariates are quite similar for households in the top half of the covariance distribution and those in the bottom half. We carefully control for income, both parametrically and with splines, at several points in our analysis.

reports the household-weighted distribution of house price covariances in our data, emphasizes this point. (Estimated covariances in our data are standardized to have zero mean and a standard deviation of 1.) Among the highly correlated city pairs that we identified in the previous figure, Boston/New York also have a high covariance. However, Dallas/Houston and Chicago/Seattle have approximately average covariances. Similarly, the negatively-correlated city pairs of Dallas/New York and Chicago/Dallas both have strongly negative covariances despite Dallas/New York having a much more negative correlation. The highest covariances in our data are among cities in the Northeast and among cities in California.¹⁷

The determinants of cross-MSA house price covariance are outside the scope of our study. That said, it is interesting to examine the relationship between covariance and other MSA-pair characteristics. Regression results (not shown) indicate that house price covariance is increasing in the covariance in per capita income, which reflects the strength of the correlation in business cycles across an MSA pair. It is also increasing in the fraction of migration flows shared between a pair. Finally, house price covariance has a nonlinear relationship with distance: Covariance initially declines sharply with distance but rises again at distances of about 3,000 miles, which reflects the strong covariances between high-priced cities on the coasts.

4 Estimation Strategy

Equation 6 reflects the notion that when a household faces particular realizations of house price growth in their origin and destination housing markets, the effect on housing consumption in the destination is dampened when house prices in the two markets move together. The hedging effect implies that if households were to draw repeatedly from the distributions of house prices in their origin and destination markets, higher covariance (more hedged) households would experience lower variance of subsequent housing consumption. We seek to estimate this relationship between covariance of house prices and the variance of subsequent housing spending, conditional on the distributions of house prices.

¹⁷Our results are robust to excluding regional groups of households, such as all those who moved from or to an MSA in California.

One empirical approach would be simply to test whether the the change in housing consumption between the purchase of one house in an origin city and another in the destination differs depending on whether the house price growth over that same time period was similar in the origin and destination housing markets or not. That approach would parallel Gruber’s (1997) study of the consumption-smoothing effects of unemployment insurance, which estimated whether the change in consumption around an unemployment shock varied with the generosity of unemployment insurance. However, unlike in other consumption-smoothing research, we do not directly observe the shock — a household’s realization of house price growth.

Instead, our empirical strategy adapts the within-household prediction of equation 6 to our cross-sectional data. Equation 6 refers to a variance in housing consumption taken over a set of possible realizations of house prices for one representative household. Since we observe only one housing choice per household, we estimate the variance of housing consumption across a number of households, each having drawn one realization from a common distribution, and relate that to the covariance in house prices between the origin and destination.¹⁸

As the example in the introduction emphasized, we are in essence comparing households who moved between high-covariance city pairs such as Boston and New York to households who moved between low-covariance city pairs such as Dallas and New York during a five-year window. Each of these households receives a different set of origin and destination house price shocks depending on which city pair they move between and when during the 1995 to 2000 period they move. We implicitly pool households by the covariance of their city pairs, using cross-sectional differences in house price realizations across city pairs to trace out the conditional variance of household-level housing consumption. Additional identifying variation comes from households having different price realizations due to moving at different times during the five-year window. The key identifying assumption is that, conditional on the covariates, households draw from the same distribution of house price shocks. We take several steps to make this assumption palatable.

First, we calculate the mean of house values conditional on demographic characteristics, in

¹⁸Although common in time-series and financial econometrics, conditional variance estimation for its own sake is fairly rare in cross-sectional applications. See Shore (2010) for one example. Carroll and Ruppert (1988) provides a useful summary of the literature on estimating conditional variance functions to that date.

essence estimating the deviation of realized housing spending from predicted spending. We also condition on fixed effects for the origin and destination MSAs, accounting for differences among them that could affect housing wealth or spending, such as the level of house prices. Second, when we relate the variance of house values to the covariance in house prices, we again condition on household characteristics and origin and destination MSA fixed effects. This controls for the possibility that the demographic composition of movers across MSA pairs could vary in a way that is systematically related to the covariance. It also controls for all households departing an origin MSA having systemically higher variance in wealth, or a destination MSA imposing more variance in housing values on movers, independent of the MSA of origin.

Finally, we can use variation at the household level to relax the assumption of the same (conditional) wealth distribution for all movers from an origin city and instead allow that distribution to vary by MSA pairs. Equation 6 predicts that, among households who move between a origin-destination city pair, the effect of covariance on the conditional variance in subsequent housing spending should be largest for households who owned a large quantity of housing before moving. We proxy for the quantity of housing owned with whether a household has a low income relative to house prices in the origin MSA. We also estimate a regression model to fit the quantity of housing owned conditional on a large set of covariates and fixed effects. To test this prediction, we interact these two measures with covariance and estimate the conditional variance with origin and destination fixed effects, as well as origin x destination fixed effects, at the cost of losing our estimate of the main effect of covariance. This allows us to look *within* the group of, for example, New York-Boston movers, holding constant all the factors specific to that pair of cities.

Our estimation strategy also addresses another empirical challenge that follows from the fact that households infrequently adjust their housing consumption. Prior consumption-smoothing research has used households' pre-shock consumption as a measure of their desired consumption. Since households rarely adjust their housing consumption, we cannot easily compute the change in housing consumption by taking the difference across two years. For most households that difference is zero and, even for households who move, both their prior housing consumption and their new consumption can be quite different from the latent desired amount of housing consumption assumed

in the theory (Edin and Englund 1991). In any case, in our data we also do not observe housing consumption prior to moving.¹⁹

Using the same notation as in Section 2, we assume that our data are generated by the following heteroskedastic model, where $q_{i,t}$ is the value of the house purchased by household i when it moves from $m_{i,t-1}$ to $m_{i,t}$ at time t . In addition, X_i is a vector of covariates other than covariance:

$$P_{2000}^{m_{i,t}} q_{i,t} = \alpha Cov [P_t^{m_{i,t-1}}, P_t^{m_{i,t}}] + X_i \beta + \eta_i \sqrt{e^{\delta Cov [P_t^{m_{i,t-1}}, P_t^{m_{i,t}}] + X_i \gamma}} \quad (7)$$

$$E [\eta_i | X_i] = 0$$

$$Var [\eta_i | X_i] = 1$$

The specified functional form for the variance is convenient because it both guarantees positive predicted variances and allows us to interpret changes in variance in (approximate) percentage terms.²⁰ We include the additional covariates to control for cross-sectional differences in household attributes that are correlated with covariance and might affect the quantity of housing purchased.

Equation 7 emphasizes that self-reported house values, which are what is reported in the Census data, are a measure of the price per unit quantity of housing in 2000 ($P_{2000}^{m_{i,t}}$) multiplied by the quantity purchased when the household moved in year t , whereas the theoretical results in Section 2 relate only to the quantity of housing consumed. However, since all households in the Census report their house values at approximately the same time, we can assume that the price per unit quantity is the same for all agents in a given destination MSA. Then, when we take the log of the squared residuals in the second stage, the fixed effect for the destination MSA absorbs the MSA-level price component in the variance term. The remaining differences in the log variance of house values must reflect differences in the variance of the quantity owned.²¹

¹⁹Housing tenure choice in the origin also is not available in the IPUMS. Among homeowners in the Panel Study of Income Dynamics (PSID) who moved across state lines in the prior year, about 60 percent were previously homeowners. Since households who were renters in the origin would not benefit from a hedge, this data omission should make it more difficult for us to find an effect of covariance on the mean and variance of housing consumption.

²⁰We examine the functional form further below by estimating a nonlinear model using splines.

²¹Our baseline specification does not allow for interactions between destination house price levels and covariance or the other covariates (X_i) to affect *mean* housing expenditures. We discuss this point further in the results section below.

We estimate the conditional variance of house values in two stages. In the first stage, we regress house values ($P_{2000}^{m_{i,t}} q_{i,t}$) on origin-destination covariance ($Cov [P_t^{m_{i,t-1}}, P_t^{m_{i,t}}]$), full sets of origin and destination fixed effects, and the vector of covariates X_i from equation 7.

$$P_{2000}^{m_{i,t}} q_{i,t} = \alpha Cov [P_t^{m_{i,t-1}}, P_t^{m_{i,t}}] + X_i \beta + \lambda_{m_{i,t-1}} + \lambda_{m_{i,t}} + \eta_i \quad (8)$$

This regression yields an estimate of the conditional mean of house values as well as conditionally mean-zero residuals ($\hat{\eta}_i$). We then run the second-stage regression

$$\log(\hat{\eta}_i^2) = \delta Cov [P_t^{m_{i,t-1}}, P_t^{m_{i,t}}] + X_i \gamma + \rho_{m_{i,t-1}} + \rho_{m_{i,t}} + \nu_i \quad (9)$$

where ν_i is an error term defined by $E[\nu_i | Cov [P_t^{m_{i,t-1}}, P_t^{m_{i,t}}], X_i, \rho_{m_{i,t-1}}, \rho_{m_{i,t}}] = 0$. Equation 9 follows from the facts that (1) we can always write a variable as the sum of its conditional expectation plus a conditionally mean-zero error term and (2) we can “plug in” consistent estimates of the first-stage errors (i.e., the residuals) on the left side and still get consistent estimates of δ and γ from equation 7.²² We report bootstrapped standard errors for the conditional variance (second stage) estimates.

4.1 Selection Bias

Our model assumes that all households own their homes and have no control over their migration decisions. A potential critique of our empirical approach is thus that households might change where they choose to move, or whether they buy or rent, based on how well the price of the houses they sold tracked prices in the MSA they moved to. However, it turns out that any potential bias would make it harder for us to discern an effect.

We can reasonably reject any concern about differential migration — choosing a destination because of covariance — because it requires that a substantial fraction of households change their migration decisions in response to price swings in the destination or origin. In practice, migration

²²Note that using squared residuals biases the estimate of the conditional variance function, although it remains consistent. This can be corrected by studentization of the residuals (Carroll and Ruppert 1988, p. 78). The correction has virtually no effect on our estimates, so we leave it out to maintain simplicity.

flows are nearly constant from year to year. Using U.S. Internal Revenue Service data on migration from the 1980s to the present, we regressed the logarithm of the number of households moving across an MSA pair in a given year on MSA-pair dummies. The R^2 from this regression is about 0.95, leaving only a small set of households that could be affected by annual swings in house prices. If we add to the regression the percentage by which the destination house price is above or below the origin house price, we get an elasticity of about 0.1, which is extremely precisely estimated thanks to the large sample size. Even if we multiply this elasticity by 0.2, which is about the 99th percentile of absolute annual swings in the house price gap, we still get a change in the number of migrants across a given city pair of only 2 percent. An average house price swing would shift migration by less than 0.5 percent.

A second potential issue could arise if households who experienced low covariances are less likely to buy a house after moving and thus do not show up in our sample of home owners.²³ This pattern is precisely what we find in section 5.4 below. That said, any selection bias from the endogeneity of the tenure decision most likely leads us to underestimate the effect of covariance on the variance in housing spending.

Consider again the empirical model in Equation 7 that relates house value to cross-market house price covariance. Let $H_{i,t}$ be an indicator variable for whether household i purchases a home after moving to location $m_{i,t}$. We observe this household's desired housing purchase only if $H_{i,t} = 1$. Suppose that $Pr[H_{i,t} = 1]$ depends on the same set of covariates, including fixed effects, as $q_{i,t}$, so that

$$H_{i,t} = 1 [\alpha' Cov [P_t^{m_{i,t-1}}, P_t^{m_{i,t}}] + X_i \beta' + \phi_i > 0]$$

$$corr(\eta_i, \phi_i) \neq 0$$

This is similar to a standard selection model, where the conditional probability of owning and conditional housing consumption can be modeled as correlated random variables.²⁴

²³This would require the price to deviate from the present value of renting. One possibility is that the user cost relationship is not constant. Alternatively, the household may have a short horizon, as in Sinai and Souleles (2005) or Campbell and Cocco (2007).

²⁴To actually estimate a selection-corrected version of the model that is not identified solely from functional form assumptions would require an instrument that affects the probability of owning but not the demand for housing.

Assuming that the correlation is positive — so households who have an unobservable taste for home owning also desire relatively more expensive houses — the conditional distribution of house prices is probabilistically truncated from the left. That is, households from the low end of the house price distribution are more likely to opt in and become homeowners when the covariance rises. This biases estimates of both the conditional mean and conditional variance.

If η_i is conditionally normally distributed, then the lesser probabilistic truncation as covariance rises leads to an *increase* in the conditional variance of house values. This result is straightforward to show for the case in which η_i and ϕ_i are conditionally bivariate normal, although it does not hold for all distributions in general. Since the distribution of home values in our sample, conditional on covariates, does appear to be approximately normal, any selection should bias our estimates toward zero.

5 Results

5.1 Conditional Variance Estimates

We present our results in two parts. The first section relates the conditional variance of house values to the covariance between origin and destination MSA pairs, while the second examines the possibility that the covariance hedge might also make it easier for households to buy a house, rather than rent, in the destination.

In Table 3, we report our main result: The conditional variance of house values among households who move between more highly covarying MSAs is lower, even controlling for a wide set of covariates in both the housing demand stage and the variance stage. The estimates reported in this table are from the second stage of the conditional variance estimation laid out in Section 4, which relates the conditional variance in house values across households who move between an MSA pair to the covariance in house prices of the two MSAs.

In the first column of Table 3, we estimate the log variance of house prices conditional on origin-destination covariance and full sets of origin and destination dummy variables, but no household-level covariates. We standardize covariance so that the estimated coefficient can be interpreted as

the marginal effect of a one-standard deviation change. Our estimate in column (1) is that a one standard-deviation increase in cross-market covariance reduces the variance of destination house values by about 17 percent, with a standard error of 2.4 percent.

By including origin and destination dummies, we control for any origin- or destination-specific differences in the variance of house values that might be correlated with covariance. In essence, we are comparing households who move to MSA B from an MSA that does not covary with it to those who move to MSA B from an MSA that does covary with it. Any factors that are specific to B are absorbed by the destination fixed effects. Likewise, we simultaneously compare households who move from MSA A to either high or low covariance destinations. Factors specific to A are absorbed by the origin fixed effects.

For example, if people who move from San Francisco have more variable (unobserved) wealth, it will be picked up by the San Francisco dummy and will not contaminate our estimates. Similarly, if the New York MSA happens to have a wider variety of house values than other MSAs, the New York fixed effect will absorb that. Instead, we rely on households who move from San Francisco to more highly covarying MSAs having lower variance of house values at their destinations than other households who moved out of San Francisco, and households who move to New York from more highly covarying markets having a lower variance of house values in New York than other households who moved to New York.

The same covariates in the second stage are also included in the first stage regression. The origin and destination fixed effects pick up differences in the mean house value across MSAs as well as average differences in housing demand among movers to and from each MSA. Meanwhile, the covariance term in the first stage picks up differences in the conditional *mean* that are correlated with covariance.

Column (2) repeats the conditional variance estimation, this time adding controls for household characteristics such as family size and the age, sex, citizenship status, race, education, and marital status of the household head. These covariates serve two purposes. In the first stage, the estimation of the conditional mean, the covariates control for differences in latent housing demand that are functions of observable household characteristics, the characteristics of the origin and destination

cities, and the effect of covariance. In the second stage, the covariates control for differences in the composition of movers that might be correlated with the variance of housing demand. For example, if highly-educated households had more or less variability in housing demand and were more or less likely to move between covarying MSAs, our estimate would be biased in the absence of the controls. With the addition of the household controls, the estimated covariance coefficient shrinks to -0.127, with a standard error of 0.019.

In column (3) we add current household income and its square to the set of covariates, as proxies for lifetime income. While current income is probably endogenously determined by households, we include it to make sure that our estimated relationship between covariance and variance of house values does not simply reflect sorting of households of different incomes into different locations. This is particularly important given that the distribution of households in the top and bottom halves of the covariance distribution appear to be similar on all observable dimensions except for income. Once income is included, the coefficient of interest shrinks in magnitude, to about 10 percent, but remains quite precisely estimated. We also find that the volatility of housing demand increases with current income for most households, although the quadratic term captures the fact that it eventually declines at very high levels of income.

In column (4), we attempt to isolate the direct hedging effect of covariance. As discussed in Section 2.4, one potential confounding factor is that households may alter their initial consumption and investment choices in response to their anticipated covariances. That can induce an independent, second-order effect on the amount or volatility of wealth a household faces when purchasing a home after a move. It is worth emphasizing that this is not a statistical bias per se: Our approach does estimate the effect of covariance on the variance of housing demand. Rather, in column (4) we are interested in seeing how much of the effect is due directly to the hedging property of home owning versus the other channels by which covariance might operate.

To decompose these mechanisms, we make use of the notion that pre-move housing demand and savings decisions are based on an *ex ante* expected covariance of house prices between the origin city and all MSAs the household might move to. But *ex post* housing demand, after the move, depends only on the covariance of house prices between the origin MSA and the MSA the household

ended up moving to. By including the *ex ante* expected covariance as a covariate, we can control for the effect of the non-hedge channels on housing demand and the volatility of wealth, while still identifying the hedging mechanism through our usual covariance variable.

Following Sinai and Souleles (2009), we compute the expected covariance for each household as the weighted average covariance from the origin MSA to all other MSAs, where the weights are the imputed probability that household moves between each MSA pair. Using their city of origin and the industry of employment of the household head, we calculate the household's probability of moving to each possible destination city as the rate of moving in each MSA pair x industry cell in the IPUMS. We then construct a weighted covariance using these probabilities as weights, under the assumption that households expectation of the probability of moving to a given city is the same as the actual probability for their origin-industry cell.²⁵ In column (4), we see that controlling for the indirect mechanisms barely changes our estimated effect of covariance (to -0.110) and that the estimated coefficient on expected covariance is not statistically distinguishable from zero at any conventional level of significance.

To test the robustness of our estimates, we have tried including a large set of additional covariates, both individually and jointly, as well as altering our estimation sample to exclude different groups. For example, we controlled for city-pair characteristics such as the geographic distance between the two MSAs in the pair, the amount of migration between them, the covariance of their per capita incomes, whether they are in the same state, and their industry or occupational similarity. We have also tried excluding high-price MSAs or different regional groups of MSAs, such as those in California or those in the Northeast. In none of these cases do we find a substantially different result for the main covariance effect.

Our baseline conditional variance specification using the logarithm of the squared residuals does not allow for interactions between destination house price levels and covariance or the other covariates to affect mean housing expenditures. We extended our empirics in two ways to address this possibility: First, we estimated the effect of covariance on the log variance of *log* housing

²⁵Separating the expected covariance from the realized covariance relies on the assumption that households' true expected covariances are better proxied by our weighted average measure than the covariance they realize across their actual origin and destination.

expenditures, in which case the multiplicative effect of house prices would be absorbed by the first-stage destination fixed effects. The results were qualitatively the same, although the magnitudes differed modestly due to the transformation of the left-hand-side variable (the variance of log spending rather than the variance of spending). Alternatively, in the first stage we interacted our destination fixed effects with the covariates. When feasible, this approach yielded results that were very similar to our primary results. However, because many of our covariates are categorical, the number of fixed effects in the model grows very rapidly and our ability to fully saturate the model in all of our specifications was limited by computational power.

5.2 Nonlinear Estimates

Estimating a linear model of the effect of covariance in house prices on the variance of housing spending masks the fact that there is a larger effect at high covariances than at low covariances. We estimate the nonlinear relationship using a generalized additive model (GAM) that specifies that the conditional mean of the dependent variable comprises the sum of a set of nonlinear functions, one for each covariate. The procedure estimates splines that “penalize” the likelihood function for additional degrees of freedom; this helps to avoid over-fitting (Hastie and Tibshirani 1990, Wood 2006). Wood’s (2006) recent technical innovation allows for the estimation of a GAM using automated cross validation methods to choose the penalty parameters for the spline.

We estimate a model that includes age and income as continuous covariates, using splines, as well as the usual fixed effects from the models above.²⁶ As in our more parametric versions, we estimate the model in two steps, first fitting the mean as a function of these covariates and then running another GAM with the log squared residuals on the left side and the same covariates on the right. The curve relating covariance and the variance of post-move housing spending, along with a 95 percent confidence interval, is shown in Figure 3. For covariances below the mean, which is standardized to zero, the curve is less steeply sloped than for those above the mean. On average, the slope above the mean is about -0.14, while the slope below it is just -0.06. Because of the long right tail of the distribution, about two-thirds of the sample have covariances below the mean.

²⁶To focus the plots on the dense part of the covariance distribution, we drop the top and bottom 1 percent of our sample by covariance. This has little effect on the estimated curves, precisely because they are estimated flexibly.

Households with a covariance in the 99th percentile of our sample have 40 percent less variance in destination housing spending than comparable households with little or no covariance between origin and destination. This nonlinearity is not surprising since our theoretical result, Equation 6, prescribes a linear effect only through the use of a Taylor approximation.

5.3 Within City Pair Identification

The identification in Table 3 requires that there are no unobservable differences in the variance of housing demand among movers between MSA pairs that happen to be correlated with differences in the covariance in house prices between the MSAs in the pair, conditional on origin and destination MSA fixed effects. We can relax that requirement by making use of equation 6, which predicts that covariance should decrease volatility by the most when the household owned more housing in the origin.

Empirically, we can compare the variance among movers who were likely to own large houses to the variance among those who were likely to own smaller houses to see if covariance has a larger dampening effect for the former group. A simple way to test for this effect is to follow Equation 6 literally and interact covariance with factors that shift q_{t-1} . We try two strategies: First, we interact covariance with the (standardized) ratio of household income to origin median house price, since higher-income households or those in lower-priced areas should own more housing, all else equal.²⁷ Second, we run a regression of log house values on a set of household covariates and location fixed effects and predict the quantity of housing a household was expected to own in the origin given its particular covariate values and fixed effect. We likewise interact this predicted housing consumption variable with covariance and include it in the conditional variance estimates.

One caveat to this approach is that households with different income or wealth should also have different marginal elasticities of demand ($q_P(\bar{w}, \bar{P})$ and $q_w(\bar{w}, \bar{P})$) in Equation 6. For example, richer households should be less impacted by a one-dollar increase in house prices. Differences in these elasticities would also lead to different effects of covariance, since they multiply the covariance term. That said, we expect any changes in marginal demand elasticities to be second-order relative

²⁷In principle we would like to use pre-move incomes or wealth, but these are not reported in our data. Consequently, we use post-move income as a proxy for pre-move income.

to changes in actual demand.

In columns (1) and (2) of Table 4, we find that higher covariance in house prices between an MSA pair reduces the variance in house values more for households who have high incomes relative to origin house prices or high predicted q_{t-1} . In column (1), we see that a one standard deviation increase in income increases the effect of covariance by about half relative to the mean, with a coefficient of -0.076 on the interaction term. The controls include the full set of household demographics, origin and destination MSA fixed effects, and our proxy for the *ex ante* expected covariance, paralleling column (4) of Table 3. The interaction is highly statistically significant, since the standard error is just 0.012. Similarly, in column (2) a one standard deviation increase in predicted house size substantially and statistically significantly amplifies the effect of covariance, by 4.1 percentage points. These results match our predictions based on Equation 6.

Since both income and predicted q_t vary at the household level, we can also include MSA-pair fixed effects to control for unobservable differences across MSA pairs and test whether covariance has a larger dampening effect for high q_{t-1} movers within a given MSA pair. This comes at the cost of not being able to estimate the main effect of covariance, since it varies only by MSA pair. The estimated effects in columns (3) and (4) are smaller than in (1) and (2), respectively. The coefficient on the interaction with income declines by about half (-0.029) but remains statistically significant at the 5 percent level. In column (4), we find only a very small negative effect of predicted house size on the effect of covariance on housing consumption variance; this coefficient is appreciably smaller than its standard error. Nonetheless, taken as a whole the evidence supports the idea that households who we expect had larger houses before moving have bigger reductions in variance in the destination.

5.4 Ex Post Probability of Owning

Covariance could also affect households' tenure choice in the destination, as discussed in Section 4.1. In this subsection, we test the tenure choice/covariance relationship by estimating a linear probability model, with an indicator for whether a household owns or rents its house as the de-

pendent variable, and the same sets of covariates as in the conditional variance estimates.²⁸ The estimation sample contains approximately twice as many households as our sample of homeowners.

The baseline results are reported in Table 5. In column (1) the only controls are origin and destination fixed effects. We find that households who faced higher covariances — that is, had a better hedge against destination house prices — are substantially more likely to own their homes after moving. A one standard deviation increase in covariance raises the probability of owning by about 3.0 percentage points (with a standard error of 0.3 percent). In our sample, the average home-ownership rate is 50 percent, so 3 percentage points corresponds to a 6 percent increase.²⁹ Columns (2) through (4) sequentially add the full vector of household controls that were used previously in Table 3. The estimated coefficient on covariance is not affected much, and ranges from 0.030 to 0.025.

Table 6 repeats the strategy from Table 4 of interacting covariance with factors that shift pre-move housing quantity q_{t-1} . With or without origin x destination fixed effects, we find small and statistically insignificant effects of the income-origin price ratio on the covariance hedging effect. Large predicted house sizes do have an effect, with a one standard deviation increase in predicted q_{t-1} increasing the effect of covariance by about 1 percentage point, with a standard error of .2 percentage points. All told, the covariance hedge appears to work not only by reducing the variance of subsequent housing consumption for home owners but also by increasing the probability that a household will be able to afford a home at the destination.

6 Magnitudes

In this subsection, we demonstrate the scale of the covariance hedge by computing the predicted reduction in variance across several groups and parts of the covariance distribution. We find that the hedge is strongest for households who are likely to own larger homes and households who move

²⁸We have also estimated probit models; the results are very similar. We prefer the linear probability model because fixed effect probit estimates are not necessarily consistent when the number of observations within each group are fixed. This is the “incidental parameters” problem. Another alternative candidate model, fixed effects logit, can only accommodate a single set of fixed effects, where we have several. Including additional sets of dummy variables reintroduces the issue of incidental parameters (Wooldridge 2002, pp. 491-492).

²⁹The home-ownership rate in the selected sample of inter-MSA migrants is substantially lower than the national home-ownership rate because renters are much more likely to move than homeowners.

cities at the high end of the covariance distribution. The first column of Table 7 uses our estimates to calculate the percentage effect of a one standard deviation increase in covariance on post-move housing consumption variance. In addition to the average effect across our entire sample, taken from column (4) of Table 3, we also calculate the effect for households with high or low incomes relative to prices in the origin, using column (1) of Table 4.³⁰ Households in the high-income group, which we define as having an income-origin price ratio at the 90th percentile, were more likely to own a larger home before moving and thus get a larger average benefit from higher covariance, of about 18 percent for each standard deviation. Conversely, the low-income group, at the 10th percentile of the income-origin price ratio, gets just a 7 percent reduction.

Since a standard deviation of covariance is not an especially intuitive measure, it is perhaps more useful to compare the strength of the hedge at different points in the covariance distribution. Because the distribution has a long right tail, as seen in Figure 2, similar percentile increases in covariance have a larger effect at the top of the distribution than at the bottom. For example, for the average household, moving between cities at the 5th percentile of covariance versus a city pair at the median reduces variance by just 6 percent. Moving between the 95th percentile city pair versus the median, on the other hand, reduces variance by 24 percent. Much of this effect is concentrated at the top of the distribution: A household with a covariance at the 95th percentile gets an 18 percent reduction in variance relative to a household at the 75th percentile.

The same pattern holds for the high- and low-income groups. The strongest predicted covariance hedge is for high income households at the top of the covariance distribution, who have 40 percent less variance in post-move consumption as high income households at the median of covariance. Meanwhile, low-income households in the lower half of the covariance distribution get much less benefit from the hedge. A household with low income relative to prices in their origin city who moves between cities at the median of covariance experiences just 4 percent lower variance than a household who moves between cities at the 5th percentile.

³⁰For ease of explanation, we do not use the nonlinear estimates from the GAM described above and shown in Figure 3. The patterns described here would be even more stark if we did.

7 Conclusion

In this paper, we examine the empirical link between cross-market house price covariance and variance in subsequent housing consumption. Theory suggests that higher covariance should hedge the volatility of housing consumption since changes in home owners' wealth would offset changes in the costliness of housing. Prior research has demonstrated that households respond prospectively to a *potential* hedge by being more likely to own their houses and to spend more on housing when the potential hedge is stronger. This paper shows that the hedge works.

Empirically identifying whether home ownership successfully reduces the volatility of housing consumption is challenging because adjusting housing consumption is a low-frequency event, and in our data we observe neither the shock to housing costs nor the household's latent housing demand. We surmount these difficulties by applying a conditional variance estimation technique that is novel in the consumption smoothing literature. In essence, we use the variance of housing spending across a cross-section of households as an estimate of the variance of housing spending that a single household would experience across different states of the world.

This strategy works because we can condition on household level observable and MSA-level unobservable characteristics, so each household varies only by the (unobserved) shock to housing costs. Differences in the variance of housing spending among movers to a destination can thus be related to differences faced by those households in the covariance in house prices between their origins and destinations. We also examine whether the variance of housing consumption for households who move between a given pair of MSAs responds more to covariance for those households that theory predicts would be more sensitive to it, that is, households who were likely to own larger homes (and thus have a bigger hedge) before they moved.

Our estimates show that home ownership significantly reduces the variance in housing spending for households that move between covarying MSAs. A one standard deviation increase in covariance, holding all else constant, reduces the average variance of housing spending by 10 to 17 percent, depending on the specification. This average estimate masks considerable nonlinearity and heterogeneity across groups. Allowing the estimated coefficient to vary non-linearly with the level of covariance by using a generalized additive model, we find that for households with covariance

above the mean, a one standard deviation change in covariance would reduce the variance of housing spending by 14 percent whereas households with below-mean covariance enjoy just a 6 percent reduction. The effect is especially sizable for wealthy households (20 percent) as well as those who are particularly likely to have owned a large home before moving. In addition, households who face higher covariances are more likely to purchase a house in the destination since they are better protected against unexpected changes in house prices.

Finally, we show that the hedge can be especially valuable for certain households, particularly those who own large homes and move across cities whose prices covary strongly. We find that for such households, covariance can reduce the variance of post-move housing spending by more than 40 percent relative to otherwise identical households who move across cities with covariance at the median. The variance in housing consumption is reduced even more because greater covariance raises the odds that a household can afford to own a home after a move. Conversely, the hedging benefit is weakest for low-income households, who do not own much if any housing, and households who move across city pairs that do not covary much.

This natural hedge provided by home owning can help explain some facts that the conventional wisdom finds surprising. For example, the measured marginal propensity to consume out of housing capital gains might be low, as found by Calomiris, Longhofer and Miles (2009), Attanasio, Blow, Hamilton and Leicester (2009) and Campbell and Cocco (2007), because increases in housing wealth are spent on commensurately higher housing costs. As another example, while insurance markets have arisen to mitigate most other major sources of consumption uncertainty — health care, long-term care, or even college tuition costs — markets to insure against house price uncertainty have not taken off (Shiller 2008). Our results suggest that simply owning a house provides valuable insurance against housing costs in future cities, obviating some of the need for a separate financial product. Finally, higher covariance in house prices may mitigate not only changes in housing consumption after a move, but changes in non-housing consumption as well.

A Mathematical Appendix

We first demonstrate that $q_w(c_t, q_t) = \frac{\partial q_t}{\partial w_t}$ is positive under standard assumptions, namely that housing and consumption are complementary or at least sufficiently non-substitutable and that the return on the financial asset r_{t+1} exceeds the return on housing $\pi_{t+1}^{m_t}$ in expectation. Denote the Euler equations 3 and 4 by A and B , respectively. Totally differentiating the Euler equations and applying Cramer's rule indicates that we can determine the sign of $\frac{\partial q_t}{\partial w_t}$ by examining the sign of $-A_w B_s + A_s B_w$, where A_w is the partial of the first Euler equation with respect to w_t , and so on.

$$\begin{aligned} \text{sign} \left(\frac{\partial q_t}{\partial w_t} \right) &= \text{sign}(-A_w B_s + A_s B_w) \\ &= \text{sign}(u_{cc}(c_t, q_t) P_t^{m_t} \\ &\quad \cdot \beta E [u_{cc}(c_{t+1}, q_{t+1}) ((1 + r_{t+1})(r_{t+1} - \pi_{t+1}^{m_t}))]) \\ &\quad - u_{qc}(c_t, q_t) \beta E [u_{cc}(c_{t+1}, q_{t+1}) (1 + r_{t+1})^2]) \end{aligned}$$

Since $u_{cc}(\cdot)$ is negative by assumption, the first compound term is positive so long as $E[r_{t+1} - \pi_{t+1}^{m_t}] > 0$, which must be true under a no-arbitrage condition (assuming the returns on the two assets are adjusted for differences in risk) since part of the return on housing is consumed in-kind as a service flow. The second term must simply be sufficiently non-negative so as not to exceed the first term, which will be true as long as u_{qc} is positive — housing and consumption are complementary — or sufficiently non-negative.

Next, we show that $q_P(c_t, q_t) = \frac{\partial q_t}{\partial P_t^{m_t}}$ is negative, holding future prices constant.³¹ Following the same procedure as above, we totally differentiate the Euler equations and apply Cramer's rule

³¹It is natural to hold future prices constant because a transitory increase in price today should be accompanied by a decrease in the expected return.

to yield

$$\begin{aligned}
\text{sign} \left(\frac{\partial q_t}{\partial P_t^{m_t}} \right) &= \text{sign} (-A_P B_s + A_s B_P) \\
&= \text{sign} (-u_{cc} (c_t, q_t) q_t P_t^{m_t} \\
&\quad \cdot \beta E [u_{cc} (c_{t+1}, q_{t+1}) ((1 + r_{t+1}) (r_{t+1} - \pi_{t+1}^{m_t}))]) \\
&\quad + u_{qc} (c_t, q_t) q_t \beta E [u_{cc} (c_{t+1}, q_{t+1}) (1 + r_{t+1})^2] \\
&\quad + u_c (c_t, q_t) (u_{cc} (c_t, q_t) + \beta E [u_{cc} (c_{t+1}, q_{t+1}) (1 + r_{t+1})^2]))
\end{aligned}$$

This expression is quite similar to the previous one, except it has the opposite sign and includes an additional term that incorporates $u_c (c_t, q_t)$, the first derivative of the utility function with respect to consumption. This final term is unambiguously negative, so the requirements for housing demand to be decreasing in current prices are even weaker than for it to be increasing in wealth. This is not surprising, since it would be quite strange if the model implied that housing were a Giffen good.

Finally, we derive the conditions under which $\frac{\partial q_t}{\partial Cov[P_{t+1}^{m_t}, P_{t+1}^{m_{t+1}}]}$ is positive or negative. This is key to understanding how pre-move housing purchases endogenously respond to expected covariance between pre- and post-move house prices. Again totally differentiating the Euler equations and applying Cramer's rule, we get

$$\text{sign} \left(\frac{\partial q_t}{\partial Cov [P_{t+1}^{m_t}, P_{t+1}^{m_{t+1}}]} \right) = \text{sign} (-A_{Cov} B_s + A_s B_{Cov}) \quad (10)$$

Since $B_s = u_{cc} (c_{t+1}, q_{t+1}) + \beta E [u_{cc} (c_{t+1}, q_{t+1}) (1 + r_{t+1})^2]$ is negative, the sign of the first term in Equation 10 depends on the sign of A_{Cov} , the partial derivative of the first Euler equation with respect to covariance. This in turn depends on the cross-partial of $u_c (c_{t+1}, q_{t+1}) P_{t+1}^{m_t}$ with respect to pre- and post-move house prices; if this cross-partial is positive, than the expectation of the expression is increasing in covariance.

$$\frac{\partial^2 u_c (c_{t+1}, q_{t+1}) P_{t+1}^{m_t}}{\partial P_t^{m_t} \partial P_{t+1}^{m_t}} = -u_{cc} (c_{t+1}, q_{t+1}) q_{t+1} - u_{ccc} (c_{t+1}, q_{t+1}) q_t q_{t+1}$$

The first term here is positive, while the second is negative under our assumptions. This illustrates

the tension between the direct effect of covariance on the riskiness of housing investment and the indirect effect via the precautionary saving motive. If the direct effect dominates, then A_{Cov} is positive and the first term on the right side in Equation 10 is positive.

Turning to the second term on the right side of Equation 10, $A_s B_{Cov}$, we can see that A_s is negative as long as housing and consumption are not too substitutable. So the sign of this expression depends on B_{Cov} , which is negative as long as

$$\frac{\partial^2 u_c(c_{t+1}, q_{t+1})}{\partial P_t^{m_t} \partial P_{t+1}^{m_{t+1}}} = -u_{ccc}(c_{t+1}, q_{t+1}) q_t q_{t+1}$$

is negative, which it is if the marginal utility of consumption is convex, as assumed. This implies that the expectation of the expression is decreasing in covariance, so that B_{Cov} is negative. Importantly, this tends to offset the precautionary saving channel from the first expression of the right side of Equation 10, which confirms that the direct effect of higher covariance making housing less risky is likely to dominate. This means pre-move housing investment is increasing in covariance, which is in accord with the findings of Sinai and Souleles (2009), Han (2008), and Han (2010).

References

- Attanasio, Orazio P., Laura Blow, Robert Hamilton, and Andrew Leicester**, “Booms and Busts: Consumption, House Prices and Expectations,” *Economica*, 2009, 76 (301), 20–50.
- Borsch-Supan, Axel**, “Panel Data Analysis of Housing Choices,” *Regional Science and Urban Economics*, 1990, 20 (1), 65–82.
- Breusch, T.S. and A.R. Pagan**, “A Simple Test for Heteroskedasticity and Random Coefficient Variation,” *Econometrica*, 1978, 47 (5), 1287–1294.
- Browning, Martin and Thomas F. Crossley**, “Unemployment Insurance Benefit Levels and Consumption Changes,” *Journal of Public Economics*, 2001, 80 (1), 1–23.
- Calomiris, Charles W., Stanley D. Longhofer, and William Miles**, “The (Mythical?) Housing Wealth Effect,” *NBER Working Paper*, 2009, 15075.

- Campbell, John Y. and Joao F. Cocco**, “How Do House Prices Affect Consumption? Evidence From Micro Data,” *Journal of Monetary Economics*, 2007, 54 (3), 591–621.
- Carroll, Raymond J. and David Ruppert**, *Transformation and Weighting in Regression*, Chapman and Hall, 1988.
- Chetty, Raj and Adam Szeidl**, “Consumption Commitments and Risk Preferences,” *Quarterly Journal of Economics*, 2007, 122 (2), 831–877.
- Cocco, Joao F.**, “Portfolio Choice in the Presence of Housing,” *Review of Financial Studies*, 2004, 18 (2).
- Cochrane, John H.**, “A Simple Test of Consumption Insurance,” *Journal of Political Economy*, 1991, 99 (5), 957–976.
- Davidoff, Thomas**, “Labor Income, Housing Prices, and Homeownership,” *Journal of Urban Economics*, 2006, 59 (2), 329–333.
- , “Home Equity Commitment and Long-Term Care Insurance Demand,” *Journal of Public Economics*, 2010, 94 (1-2), 44–49.
- Edin, Per-Anders and Peter Englund**, “Moving Costs and Housing Demand: Are Recent Movers Really in Equilibrium?,” *Journal of Public Economics*, 1991, 44 (3), 299–320.
- Engle, Robert F.**, “Autoregressive Conditional Heteroskedasticity With Estimates of the Variance of United Kingdom Inflation,” *Econometrica*, 1982, 50 (4), 987–1007.
- Flavin, Majorie and Takashi Yamashita**, “Owner-Occupied Housing and the Composition of the Household Portfolio,” *American Economic Review*, 2002, 92 (1), 345–362.
- Glaeser, Edward L., Matthew E. Kahn, and Jordan Rappaport**, “Why Do the Poor Live in Cities? The Role of Public Transportation,” *Journal of Urban Economics*, 2008, 63 (1), 1–24.
- Gruber, Jonathan**, “The Consumption Smoothing Benefits of Unemployment Insurance,” *American Economic Review*, 1997, 87 (1), 192–205.

- , “Cash Welfare as a Consumption Smoothing Mechanism for Divorced Mothers,” *Journal of Public Economics*, 2000, 75 (2), 157–182.
- Han, Lu**, “Hedging House Price Risk in the Presence of Lumpy Transaction Costs,” *Journal of Urban Economics*, 2008, 64 (2), 270–287.
- , “The Effects of Price Risk on Housing Demand: Empirical Evidence from U.S. Markets,” *Review of Financial Studies*, 2010, 23 (11), 3889–3928.
- Hansen, Julia L., John P. Formby, and W. James Smith**, “Estimating the Income Elasticity of Demand for Housing: A Comparison of Traditional and Lorenz-Concentration Methodologies,” *Journal of Housing Economics*, 1998, 7 (4), 328–342.
- Hastie, T.J. and R.J. Tibshirani**, *Generalized Additive Models*, Chapman and Hall/CRC, 1990.
- Haurin, Donald R., Patric H. Hendershott, and Dongwook Kim**, “Housing Decisions of American Youth,” *Journal of Urban Economics*, 1994, 35 (1), 28–45.
- Hryshko, Dmytro, Maria Jose Luengo-Prado, and Bent E. Sorensen**, “House Prices and Risk Sharing,” 2010. mimeo.
- Hurst, Erik and Frank Stafford**, “Home is Where the Equity Is: Mortgage Refinancing and Household Consumption,” *Journal of Money, Credit and Banking*, 2004, 36 (6), 985–1014.
- Kimball, Miles S.**, “Precautionary Saving in the Small and in the Large,” *Econometrica*, 1990, 58 (1), 53–73.
- Ortalo-Magne, Francois and Andrea Prat**, “Spatial Asset Pricing: A First Step,” 2009. mimeo.
- **and Sven Rady**, “Tenure Choice and the Riskiness of Non-Housing Consumption,” *Journal of Housing Economics*, 2002, 11 (3), 266–279.
- Piazzesi, Monika, Martin Schneider, and Selale Tuzel**, “Housing, Consumption and Asset Pricing,” *Journal of Financial Economics*, 2007, 83 (3), 531–569.

- Rosen, Harvey**, “Housing Decisions and the U.S. Income Tax,” *Journal of Public Economics*, 1979, 11 (1), 1–23.
- Ruggles, Steven, Matthew Sobek, Trent Alexander, Catherine A. Fitch, Ronald Goeken, Patricia Kelly Hall, Miriam King, and Chad Ronnander**, “Integrated Public Use Microdata Series: Version 3.0,” 2004.
- Shiller, Robert J.**, *The Subprime Solution: How Today’s Global Financial Crisis Happened, and What to Do about It*, Princeton University Press, 2008.
- Shore, Stephen H.**, “For Better, For Worse: Intrahousehold Risk-Sharing over the Business Cycle,” *Review of Economics and Statistics*, 2010, 92 (3), 536–548.
- and **Todd Sinai**, “Commitment, Risk, and Consumption: Do Birds of a Feather Have Bigger Nests?,” *Review of Economics and Statistics*, 2010, 92 (2), 408–424.
- Sinai, Todd and Nicholas S. Souleles**, “Owner-Occupied Housing as a Hedge Against Rent Risk,” *Quarterly Journal of Economics*, 2005, 120 (2), 763–789.
- and — , “Can Owning a Home Hedge the Risk of Moving?,” 2009. mimeo.
- Wood, Simon N.**, *Generalized Additive Models: An Introduction with R*, Chapman and Hall/CRC, 2006.
- Wooldridge, Jeffrey M.**, *Econometric Analysis of Cross Section and Panel Data*, MIT Press, 2002.

Table 1: A Highly Stylized Example

(1)	(2)	(3)	(4)	(5)	(6)
				Elasticity parameters	
				Cobb-Douglas	Rosen (1979)
				$\epsilon_w = 1$	$\epsilon_w = 0.75$
				$\epsilon_p = -1$	$\epsilon_p = -1$
					Borsch-Supan (1990)
					$\epsilon_w = 0.1$
					$\epsilon_p = -0.3$
		State of the world			
		If house prices in city B ...house prices in city A			
		changed by:	must have changed by:	...and housing consumption in	
				period 2 would change by: ¹	
Perfect covariance	+20%	+20%	0%	-5%	-4%
	-20%	-20%	0%	+5%	+4%
Perfect negative covariance	+20%	-20%	-40%	-35%	-8%
	-20%	+20%	+40%	+35%	+8%

Notes: Households move from City A to City B. ϵ_w and ϵ_p are the wealth and destination house price elasticities of housing demand, respectively. See text for further details.

¹Relative to consumption with no price changes.

Table 2: Summary Statistics, By Sample

	All Homeowners Mean	Migrants Mean	Estimation Sample Mean Std. Dev.
House value	147,173	207,043	227,795 175,591
Household income	64,277	87,397	96,425 82,870
Age of head	53.76	45.72	44.97 13.75
Female head	0.28	0.23	0.22 0.42
Married	0.67	0.75	0.77 0.42
Family size	2.60	2.75	2.82 1.40
Black	0.07	0.06	0.06 0.24
College-educated	0.33	0.58	0.61 0.49
Covariance (levels)			324 mil. 584 mil.
Covariance (% changes)			0.0019 0.0028
Income / median origin price			0.660 0.633
N	3,330,743	146,012	100,851

Summary statistics for progressively smaller subsets of the 2000 Census PUMS. Individual-specific data are reported for the first person listed on the Census form, whom we call the “head”. “Migrants” includes only those households who reported moving across MSAs in the past five years. The “Estimation Sample” drops all households for whom we do not have covariance or other data, or who are severe outliers with respect to the reported price-income ratio. See text for further details.

Table 3: Baseline Percentage Effect of Covariance on Housing Expenditure Variance

Variable	(1)	(2)	(3)	(4)
Covariance (standardized)	-0.173 (0.024)	-0.127 (0.019)	-0.103 (0.019)	-0.110 (0.020)
Household income (\$1000's)			0.011 (0.000)	0.011 (0.000)
Household income -squared			-0.000009 (0.000001)	-0.000009 (0.000001)
Exp. Covariance (standardized)				0.023 (0.021)
Observations	100851	100851	100851	100851
R^2	0.099	0.135	0.162	0.162
Origin FE	X	X	X	X
Destination FE	X	X	X	X
Household controls		X	X	X

Standard errors are bootstrapped by origin x destination cluster using 500 replications to account for two-step estimation of conditional variance. Covariance and expected covariance are standardized to have mean zero and standard deviation one. Expected covariance is imputed based on origin and industry. Household controls include age and age-squared as well as indicator variables for sex of household head, family size, marital status, citizenship, race, English language abilities, and education.

Table 4: Interacted Percentage Effect of Covariance on Housing Expenditure Variance

Variable	(1)	(2)	(3)	(4)
Covariance (standardized)	-0.125 (0.020)	-0.134 (0.020)		
Covariance x Income / Origin Price	-0.076 (0.012)		-0.029 (0.014)	
Covariance x Predicted House Size		-0.041 (0.010)		-0.005 (0.011)
Observations	100851	100851	96350	96350
R^2	0.161	0.161	0.280	0.284
Origin FE	X	X		
Destination FE	X	X		
Orig.*Dest. FE			X	X
Household controls	X	X	X	X
Exp. Cov. control	X	X	X	X
Income controls	X	X	X	X

Standard errors are bootstrapped by origin x destination cluster using 500 replications to account for two-step estimation of conditional variance. Covariance is standardized to have mean zero and standard deviation one. Predicted house size is the fitted value from a regression of log house price on a full set of household covariates. Income-origin price ratio and predicted house size are standardized to have mean zero and standard deviation one in the interaction terms. Expected covariance is imputed based on origin and industry. Household controls include age and age-squared as well as indicator variables for sex of household head, family size, marital status, citizenship, race, English language abilities, and education. Income controls comprise a linear and a quadratic term.

Table 5: Baseline Effect of Covariance on Ex Post Probability of Owning

Variable	(1)	(2)	(3)	(4)
Covariance (standardized)	0.030 (0.003)	0.027 (0.002)	0.026 (0.002)	0.025 (0.002)
Household income (\$1000's)			0.003 (0.000)	0.003 (0.000)
Household income -squared			-0.000005 (0.000000)	-0.000005 (0.000000)
Exp. Covariance (standardized)				0.009 (0.004)
Observations	207472	207472	207472	199297
R-squared	0.057	0.267	0.298	0.297
Origin FE	X	X	X	X
Destination FE	X	X	X	X
Household controls		X	X	X

Linear probability model. Standard errors clustered at the origin x destination level. Covariance and expected covariance are standardized to have mean zero and standard deviation one. Expected covariance is imputed based on origin and industry. Household controls include age and age-squared as well as indicator variables for sex of household head, family size, marital status, citizenship, race, English language abilities, and education.

Table 6: Interacted Effect of Covariance on Ex Post Probability of Owning

Variable	(1)	(2)	(3)	(4)
Covariance (standardized)	0.025 (0.002)	0.025 (0.002)		
Covariance x Income / Origin Price	0.001 (0.002)		0.001 (0.002)	
Covariance x Predicted House Size		0.010 (0.002)		0.011 (0.002)
Observations	199297	199297	199297	199297
R^2	0.297	0.297	0.364	0.365
Origin FE	X	X		
Destination FE	X	X		
Orig.*Dest. FE			X	X
Household controls	X	X	X	X
Exp. Cov. control	X	X	X	X
Income controls	X	X	X	X

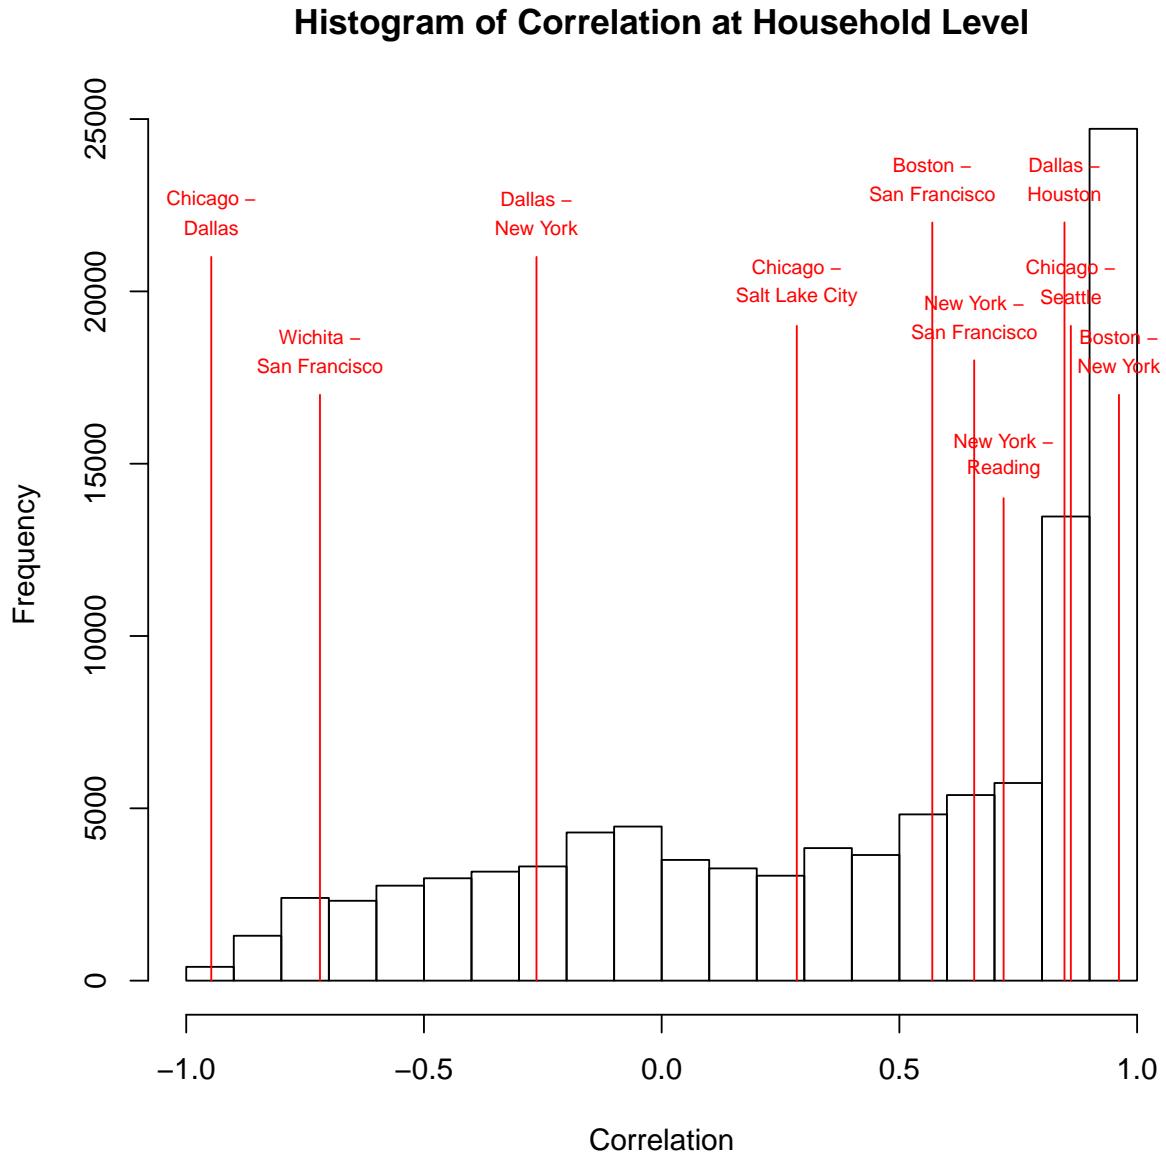
Standard errors clustered at the origin x destination level. Covariance is standardized to have mean zero and standard deviation one. Predicted house size is the fitted value from a regression of log house price on a full set of household covariates. Income-origin price ratio and predicted house size are standardized to have mean zero and standard deviation one in the interaction terms. Expected covariance is imputed based on origin and industry. Household controls include age and age-squared as well as indicator variables for sex of household head, family size, marital status, citizenship, race, English language abilities, and education. Income controls comprise a linear and a quadratic term.

Table 7: Magnitudes

	Effect of 1 SD Increase in Covariance	Effect of Shift Across Percentiles of Covariance Distribution		
		5% → 50%	50% → 95%	75% → 95%
Average Household	-10%	-6%	-24%	-18%
High Income Households	-18%	-10%	-40%	-30%
Low Income Households	-7%	-4%	-16%	-12%

We define the high- and low-income groups as households near the 90th and 10th percentiles of household income divided by origin median price, as a proxy for quantity of housing owned in the origin. The difference in the effect of covariance in the income groups comes from the income-origin price ratio interaction term in column (1) of Table 4. See text for further details.

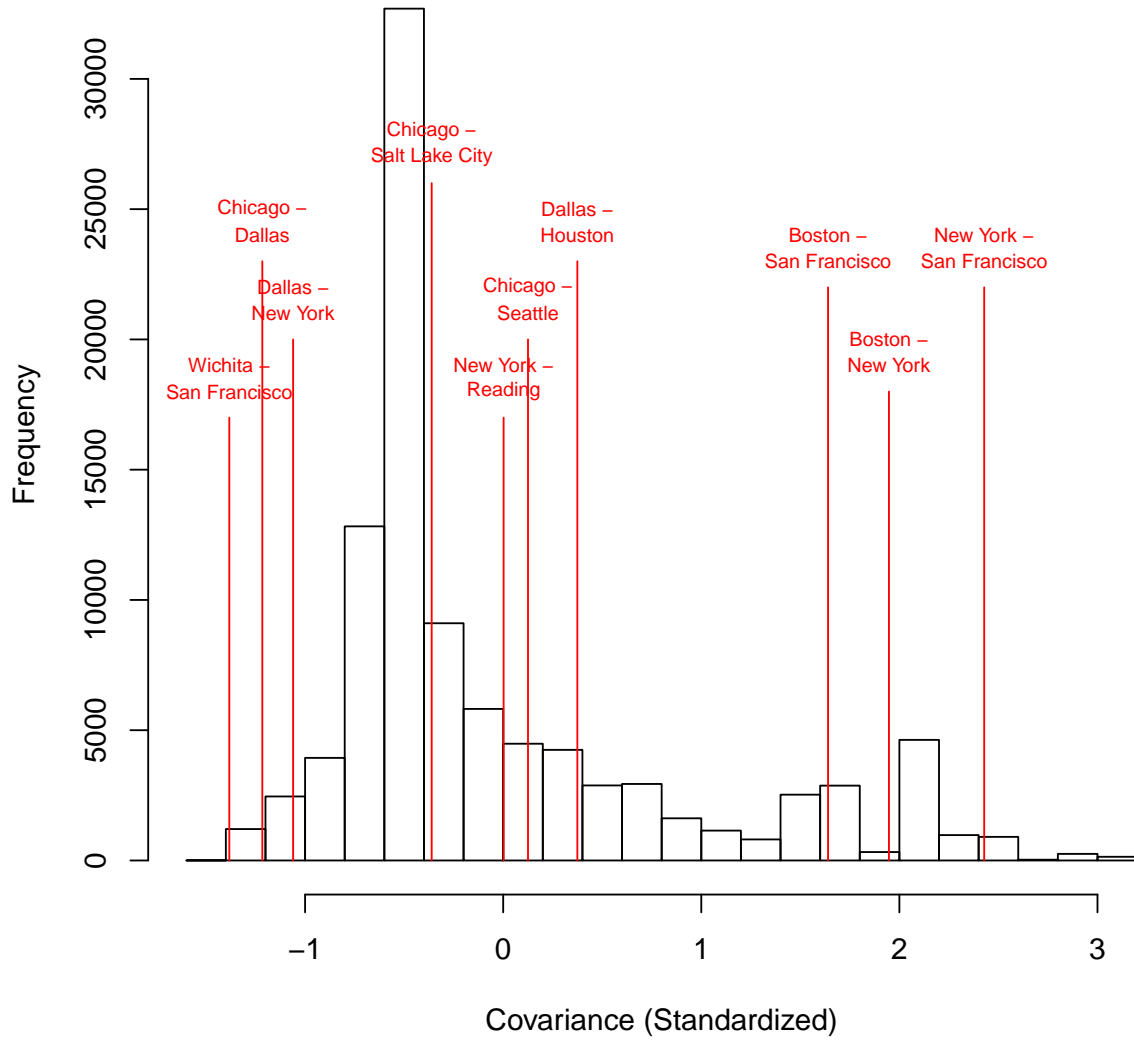
Figure 1:



Histogram of correlation at the household level, excluding the top and bottom 1 percent with respect to the covariance distribution, as in Figure 2..

Figure 2:


Histogram of Standardized Covariance at Household Level



Histogram of covariance at the household level, excluding the top and bottom 1 percent. Covariance is standardized to have mean zero and standard deviation one.

Figure 3:

Percentage Effect on Variance of House Price



Solid curve is a penalized regression spline. The regression includes the full set of discrete household covariates detailed in the text, as well as splines in income, age and (prior) expected covariance. Dashed curves show the 95% confidence interval. Covariance is standardized to have mean zero and standard deviation one.