

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

Andrew Paciorek and Todd Sinai[†]

October 22, 2010

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

*We thank Fernando Ferreira, Joe Gyourko, Chris Paciorek, David Rothschild, Albert Saiz, Stephen Shore, Nicholas Souleles, Joel Waldfogel, Justin Wolfers, Maisy Wong and seminar participants at Wharton and the Duke ERID housing dynamics conference for helpful comments and suggestions. We are grateful to the Research Sponsors Program of the Zell-Lurie Real Estate Center at Wharton for funding. All errors are our own.

[†]Real Estate Department, Wharton School, University of Pennsylvania.

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 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, or dampens the increase in housing consumption due to lower prices.

We illustrate this idea in a simple two-period representative agent model with two locations and stochastic house prices. The model predicts a negative relationship between the variance of housing consumption for households who recently moved and the covariance in the house prices of the origin and destination cities. It also predicts that the hedging benefit of high covariance would be strongest for households who own more housing.

Using household-level microdata from the U.S. Census, we find that home owners who recently moved between highly covarying cities have less conditional variance in their house values than do home owners who moved between low-covariance cities. This result remains even after controlling for origin and destination Metropolitan Statistical Area (MSA) fixed effects, so that our estimate of the hedging effect of covariance is identified solely from the pairing of MSAs. We also control for household-level determinants of housing demand.

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 specification. The reduction in variance is especially pronounced — as much as 18 percent — for home

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.

owners who 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.

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 2008b, Sinai and Souleles 2009), or has considered the hedging properties of homeownership in theory (Ortalo-Magne and Rady 2002). Other researchers have noted that, due to its nature as a consumption commitment, owning a home 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 empirical strategy relies on the notion that housing consumption for households who move between housing markets with a higher covariance should more closely match their latent housing demand since unexpected price shocks in the destination market are matched by price changes in the origin. In brief, we estimate the latent demand for housing and test whether the variance of the deviation of realized housing spending from predicted spending varies with the covariance of house prices between the origin and destination markets. The two-stage conditional variance procedure includes controls at both stages of the estimation for household demographics, the expected hedging benefit of home ownership, and origin and destination 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 covariance effect are identified from the pairing of MSAs.

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 relatively high incomes relative to prices in the origin or had high predicted housing

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.

consumption in the origin MSAs.

We find two other pieces of evidence consistent with the hedging interpretation of the relationship between covariance and the subsequent variance in housing spending. Since origin and destination house prices enter the housing demand function nonlinearly, the theory predicts that higher covariance should reduce mean housing spending after a move. We find this pattern in our data, with the strongest effect once again for households who appear to have owned larger houses prior to the move. In addition, 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 nonlinearly 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.

Our discussion proceeds as follows: In the next section, we outline a simple two-period model of housing consumption and moving. We use this model to derive the direct effect of the covariance of house price changes on the variance and mean of subsequent housing expenditures, 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 proxy for variance across states of the world. 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 and setup

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 for 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 two sets of estimates of the elasticity of housing demand with respect to wealth (ϵ_w) and house prices (ϵ_p).

In the first row of Table 1, the household faces 20 percent growth in house prices in city B, which would reduce its demand for housing there. However, in the case of perfect covariance, house prices also grow commensurately in city A, making the household 20 percent wealthier and raising its demand for housing in city B. Under Cobb-Douglas preferences with wealth and own-price demand elasticities of 1 and -1 respectively, the price effect and wealth effect exactly net out and there is no change in housing consumption after the move (in the second-to-last column).³ Under other plausible estimates of the elasticities, such as from Rosen (1979) (in the last column), 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. In this

³Glaeser and Gyourko (2007) survey estimates of price elasticities of housing demand and find -1.0 to be in the middle of the distribution.

particular example, the price elasticity dominates the wealth elasticity and housing demand falls by 5 percent. Conversely, in the second row, a 20 percent decline in house prices in city B also implies a 20 percent decline in wealth in the perfect covariance case. This yields either no net effect on housing demand in city B, or a 5 percent increase, depending on the elasticities used.

By contrast, the bottom two rows of Table 1 assume house price growth is perfectly negatively covarying in cities A and B. 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.

The intuition of the paper can be seen by comparing the positive covariance and negative covariance cases. The potential variation in housing demand in city B is much greater for the negative covariance households regardless of which elasticities we use. 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.

To generalize this intuition and provide guidance for the econometric specification, we outline a simple representative agent precautionary saving framework. Since our focus is on the relationship between house prices in current and future housing markets, our model encompasses two locations with stochastic house prices but abstracts away from other complications.⁴ We consider the decisions of an agent who lives for two periods. At the start of period 1, he lives in city *A*. In period 2 he moves to city *B*. There are no transactions costs from moving and we assume that both the timing of the move and the destination *B* are exogenously given.

In period 1, the agent receives labor income w_1 , which he must divide

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.

between non-housing consumption (c_1), housing investment (h_1), and investment in financial assets (s_1). There is no rental sector in the model, so the agent must purchase a home in the starting location with price $h_1 = p_1^A q_1$ for a house of size q_1 . For simplicity, we let one unit of housing produce one unit of housing services and define the utility function accordingly. We denote first-period preferences over numeraire consumption and housing services by $u(c_1, \frac{h_1}{p_1^A})$.

The agent receives a stochastic return of $1 + r$ per dollar invested in financial assets in the first period.⁵ Investment in housing in location A yields a return of $1 + \pi^A$. 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 in both 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.

In the second period, the agent must allocate his wealth between consumption and housing. Second-period preferences are given by $v(c_2, \frac{h_2}{p_2^B})$, which the agent maximizes subject to the budget constraint $c_2 + h_2 = s_1(1 + r) + h_1(1 + \pi^A)$. House prices in location B grow by $1 + \pi^B$, so that the second-period price per unit of housing is $p_2^B = p_1^B(1 + \pi^B)$. The agent's indirect utility function in the second period is thus

$$V(w_2, p_2^B) = \max_{h_2} v\left(w_2 - h_2, \frac{h_2}{p_2^B}\right)$$

where second-period wealth is equal to the sum of the financial and housing wealth, or $w_2 = s_1(1 + r) + h_1(1 + \pi^A)$.

Since second-period house prices and financial returns are stochastic, the agent must form (rational) expectations about the future and maximize over

⁵We do not restrict s_1 to be positive, which allows the agent to borrow to finance housing consumption. For example, if s_1 is negative and r is known with certainty, the “financial asset” 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 had correlations with stock returns of between -0.2 and 0.2.

consumption, housing, and financial investment accordingly. He discounts second-period utility by a factor of β . Hence his problem in the first period is

$$\max_{c_1, h_1, s_1} u\left(c_1, \frac{h_1}{p_1^A}\right) + \beta E [V(s_1(1+r) + h_1(1+\pi^A), p_2^B)] \quad (1)$$

subject to

$$w_1 = c_1 + h_1 + s_1$$

We make a set of standard assumptions about preferences in order to analyze the model. First, we assume that the first- and second-period utility functions $u(\cdot)$ and $v(\cdot)$ are twice continuously differentiable, strictly increasing and strictly concave in both consumption and housing. Consequently, second-period indirect utility is twice continuously differentiable, strictly increasing and strictly concave, so that the agent is risk-averse with respect to second-period wealth. Second, we assume that $V_w(\cdot)$, the derivative of the value function with respect to wealth, is twice differentiable and sufficiently convex. This implies that the agent is a precautionary saver, meaning that he saves more for second-period consumption in response to an increase in the variance of second-period wealth.⁷

Under these assumptions, the first-order conditions are necessary and sufficient for utility maximization. They are

$$-u_c(\cdot) + \frac{u_h(\cdot)}{p_1^A} + \beta E [V_w(\cdot)(1+\pi^A)] = 0 \quad (2)$$

$$-u_c(\cdot) + \beta E [V_w(\cdot)(1+r)] = 0 \quad (3)$$

where a numeric subscript denotes a time period and a letter subscript the derivative of the function with respect to that argument. As usual, these conditions reflect the trade-offs between a dollar of housing consumption today and a dollar of housing wealth tomorrow, and a dollar of non-housing consumption today and a dollar of financial wealth tomorrow.

2.2 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 first-period housing investment,

⁷See Kimball (1990) and citations therein for a full discussion of the mathematics of the precautionary saving motive.

q_1 , fixed. We consider the implications of allowing q_1 to vary in response to changes in variance or covariance in Section 2.4.

The first way that covariance affects the variance of housing demand follows immediately from the variance of the Marshallian demand for housing quantity in the second period:

$$\begin{aligned}
& Var [q_2 (w_2, p_2^B)] \\
& \approx Var [q_2 (\bar{w}_2, \bar{p}_2^B) + q_{2w} (\bar{w}_2, \bar{p}_2^B) (w_2 - \bar{w}_2) + q_{2p} (\bar{w}_2, \bar{p}_2^B) (p_2^B - \bar{p}_2^B)] \\
& = (q_{2w} (\bar{w}_2, \bar{p}_2^B))^2 Var [w_2] + (q_{2p} (\bar{w}_2, \bar{p}_2^B))^2 Var [p_2^B] \\
& \quad + 2 (q_{2w} (\bar{w}_2, \bar{p}_2^B)) (q_{2p} (\bar{w}_2, \bar{p}_2^B)) q_1 Cov [p_2^A, p_2^B]
\end{aligned} \tag{4}$$

The first (approximate) equality follows from taking a first-order Taylor approximation to $q_2(\cdot)$ at any point (\bar{w}_2, \bar{p}_2^B) . The second equality simply applies the definition of w_2 and our assumption that financial returns are uncorrelated with housing returns.

The last two lines of equation 4 shows that second-period housing demand has higher variance 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 is a normal good — demand increases in wealth and decreases in price — 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. Equation 4 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_1 term multiplying the covariance term shows that the reduction in variance should be more pronounced for households who owned more housing in the first period.

It is worth noting that greater covariance should also reduce the variance of *non*-housing consumption, thanks to the income effect: Households will have less wealth to spend on non-housing consumption as house prices rise, and vice versa. By the same intuition as that underlying Equation 4, 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 an 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.

2.3 Mean Demand for Housing

In addition to the effect on the variance of housing consumption, the covariance hedge should also have an effect on the mean, even holding first-period decisions constant, as above. To see this, we can apply a second-order Taylor approximation around the point (\bar{w}_2, \bar{p}_2^B) . After some simplifying, we have

$$\begin{aligned}
 E [q_2 (w_2, p_2^B)] &\approx q_2 (\bar{w}_2, \bar{p}_2^B) + \frac{1}{2} q_{2ww} (\bar{w}_2, \bar{p}_2^B) Var [w_2] \\
 &+ \frac{1}{2} q_{2pp} (\bar{w}_2, \bar{p}_2^B) Var [p_2^B] \\
 &+ q_{2wp} (\bar{w}_2, \bar{p}_2^B) q_1 Cov [p_2^A, p_2^B]
 \end{aligned} \tag{5}$$

The cross-partial of housing consumption with respect to wealth and price, $q_{2wp}(\cdot)$, is negative for common utility functions such as the CES.⁸ Consequently, covariance has a negative effect on mean housing spending. This result follows from the convexity of price and wealth in the demand for second period housing. Because of the convexity, the average demand for households with high variance of second period housing demand (due to low covariance) is greater than the average demand for households whose second period housing demand is more tightly distributed.

Equation 5 indicates that, after controlling for origin and destination characteristics such as expected house prices or variance in house prices as well as household characteristics, we should find empirically that a higher covariance is correlated with lower second period housing consumption. As with variance, the effect on the mean is amplified when the household owns a bigger house in the origin, all else equal.

⁸And in the case of Cobb-Douglas utility, housing spending is a constant fraction of wealth, so housing consumption is just a constant times the ratio of wealth to price.

2.4 Endogenizing Initial Housing Consumption

The analytical results above hold when average wealth entering the second period does not change in response to covariance. However, first-period housing consumption might respond to the 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 (2008b) and Han (2008a) find theoretical and suggestive 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 first-period housing consumption. One possible mechanism is that first-period housing is more valuable when the expected hedge is stronger because it is better at reducing second-period consumption volatility. It is straightforward to show the basis for this intuition using our model of housing investment.

Suppose the second-period value function $V(\cdot)$ from Equation 1 can be written

$$V(w_2, p_2^B) = f(\nu(w_2, p_2^B))$$

where $\nu(\cdot)$ is a per-period utility function and $f(\cdot)$, an increasing and concave function with a positive third derivative, reflects the household's willingness to trade wealth between periods.⁹ Applying a first-order Taylor approximation to $\nu(\cdot)$, we have

$$\begin{aligned} Var[\nu(w_2, p_2^B)] &\approx (\nu_w(\bar{w}_2, \bar{p}_2^B))^2 Var[w_2] + (\nu_p(\bar{w}_2, \bar{p}_2^B))^2 Var[p_2^B] \\ &\quad + (\nu_w(\bar{w}_2, \bar{p}_2^B)) (\nu_p(\bar{w}_2, \bar{p}_2^B)) q_1 Cov[p_2^A, p_2^B] \end{aligned} \quad (6)$$

Notice that only the last term contains covariance, and that it multiplies q_1 . Because housing is a normal good, $\nu_p(\cdot)$ is negative and $\nu_w(\cdot)$ is positive. Consequently, the variance of $\nu(\cdot)$ is decreasing in the product of q_1 and the covariance of house prices. Higher covariance amplifies the marginal decrease in $Var[\nu(w_2, p_2^B)]$ that results from increasing q_1 , and a decrease in $Var[\nu(w_2, p_2^B)]$ increases expected utility because of the concavity of $f(\cdot)$. This implies that, on the margin, higher covariance makes first-period

⁹Example functional forms that meet these criteria are Cobb-Douglas for $\nu(\cdot)$, the indirect utility function, and constant relative risk aversion (CRRA) for $f(\cdot)$.

housing q_1 more valuable.¹⁰ Putting these steps together, we have

$$\begin{aligned} \text{sign} \left(\frac{\partial q_1}{\partial \text{Cov} [p_2^A, p_2^B]} \right) &= \text{sign} \left(\frac{\partial EV (s_1(1+r) + h_1(1+\pi^A), p_2^B)}{\partial \text{Cov} [p_2^A, p_2^B]} \right) \\ &= -\text{sign} \left(\frac{\partial \text{Var} [\nu (w_2, p_2^B)]}{\partial \text{Cov} [p_2^A, p_2^B]} \right) \\ &> 0 \end{aligned}$$

Another possible mechanism for an endogenous response of first-period 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 in the the first period. As noted above, the variance of $\nu(\cdot)$ in Equation 6 is decreasing in the product of q_1 and the covariance of house prices. Since the first derivative of $f(\cdot)$ is convex, the agent is a precautionary saver. He will thus prefer to consume more resources in the first period, and fewer in the second period, when faced with a higher covariance. This suggests that the household will decrease financial saving s_1 in favor of first-period consumption and possibly the first-period housing purchase.¹¹

Regardless of the mechanism, when first-period housing consumption and saving respond to the hedging benefit of homeownership, there is a second channel — other than the direct hedging effect in Equation 4 — by which house price covariance can affect the variance of housing consumption in the second period. Namely, households who enter the second period 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 second period wealth. We further explain this approach in Section 5.

In addition, the reaction of first period housing consumption to covariance provides another mechanism, besides Equation 5, for higher covariance to affect the *mean* of second-period housing consumption. For example, households who reduce their precautionary saving due to more effective

¹⁰This heuristic argument can be made more formally through tedious but straightforward math involving a determinant and derivatives of the two first-order conditions.

¹¹Although we can derive exact conditions for the net effects of covariance on first-period choices, they depend on a large number of factors that are difficult to conceptualize.

hedgers spend less on housing in the second period. Regardless of the theoretical mechanism, any empirical relationship between covariance and subsequent mean housing spending would reject the null of no hedge. We will also try to distinguish the first-period housing consumption channel from the nonlinearity described in Section 2.3 by controlling for ex ante expected covariance.

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.¹² 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 2000.¹³ 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.¹⁴ 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.¹⁵

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 sam-

¹²Ruggles, Sobek, Alexander, Fitch, Goeken, Hall, King and Ronnander (2004)

¹³An MSA is designed to correspond to a labor market area, and typically contains one or more focal cities and their surrounding suburbs.

¹⁴The results are similar, if less well grounded in our theoretical results, if we use covariances of real house price growth.

¹⁵The 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.

ple.¹⁶ Of this group, we keep those who have moved domestically in the last five years and have both their current and previous metropolitan statistical areas of residence identified.¹⁷ This leaves a sample of about 150,000 households, across 284 origin 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 transitively 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. Figure 1 shows the distribution of covariances, standardized to have mean zero and standard deviation one, imputed to our IPUMS households. Most covariances are quite low, reflecting households who move to or from MSAs with low house price variances. However, there is a significant tail of higher covariances, due to households moving between highly correlated high-variance MSAs.¹⁸

¹⁶This group excludes the few home-owning households with a household head — the first person listed on the Census form — under the age of 25.

¹⁷Since 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 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.

¹⁸The highest covariances in our data are among cities in the Northeast and among cities in California. Our results are robust to excluding regional groups of households, such as all those who moved from or to an MSA in California.

4 Estimation Strategy

Equation 4 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. It is this relationship between covariance of house prices and the variance of subsequent housing spending, conditional on the distributions of house prices, that we seek to estimate.

One empirical approach would be simply to test whether 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. Using the variance in housing consumption will enable us to estimate the difference in response for high- and low-covariance households while controlling nonparametrically for the unobserved distribution of house price shocks.

Our approach is to estimate the conditional variance across households in housing spending in a destination city subsequent to a move and relate that to the covariance in house prices between the origin and destination.¹⁹ The key identifying assumption is that all households departing the same origin city draw from the same distribution of initial house price shocks. Equation 4 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. We take several steps to make this identifying assumption palatable.

¹⁹Although 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.

First, we calculate the mean of house values conditional on demographic characteristics, in 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 4 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.

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.²⁰ Instead, our

²⁰Housing tenure choice in the origin also is not available in the IPUMS. Among home-

estimation procedure embeds a prediction of a household’s latent demand for housing.

We assume that our data are generated by the following heteroskedastic model, where y_h is the value of the house owned by household h and X_h is a vector of covariates:

$$y_h = X_h\beta + \eta_h\sqrt{e^{X_h\gamma}} \quad (7)$$

$$E[\eta_h | X_h] = 0$$

$$Var[\eta_h | X_h] = 1$$

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.²¹

We estimate the conditional variance of house values in two stages. In the first stage, we regress house values (y_h) on origin-destination covariance (c_{ij}), full sets of origin and destination dummy variables, and a vector of household characteristics \tilde{x}_h , which together comprise the full vector of covariates X_h from equation 7.

$$y_h = \beta_0 + \beta_1 c_{ij} + \delta_i + \lambda_j + \tilde{x}_h' \beta_2 + \epsilon_h \quad (8)$$

Here h indexes households, while i and j denote the origin and destination of that household, so that $i = i(h)$ and $j = j(h)$.

This regression yields an estimate of the conditional mean of house values, $\hat{E}[y_h | c_{ij}, \delta_i, \lambda_j, \tilde{x}_h]$, as well as conditionally mean-zero residuals ($\hat{\epsilon}_h$). We then run the second-stage regression

$$\log(\hat{\epsilon}_h^2) = \gamma_0 + \gamma_1 c_{ij} + \delta_i + \lambda_j + \tilde{x}_h' \gamma_2 + \nu_h \quad (9)$$

where ν_h is an error term defined by $E[\nu_h | X_h] = 0$. This equation 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 estimate of the first-stage errors (i.e., the residuals)

owners 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.

²¹We examine the functional form further and estimate a nonlinear model using splines below.

on the left side and still get consistent estimates of γ from equation 7.²² We report bootstrapped standard errors for the conditional variance (second stage) estimates.

It is important to recognize that self-reported house values, which are what is reported in the Census data, are a measure of the current price per unit quantity of housing multiplied by the quantity owned, 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 effects for the destination MSA absorb the MSA-level price component. The remaining differences in the log variance of house values must reflect differences in the variance of the quantity owned.

We are also interested in estimating the effect of covariance on the conditional *mean* of housing consumption. Although the first stage of the conditional variance estimation routine is a model of the conditional mean of housing spending, we can only separate quantity from price if we use log spending instead of the spending level. Therefore, when discussing the conditional mean estimates, we run separate regressions of log house values on the relevant covariates. For these regressions we report standard errors corrected for heteroskedasticity and clustering at the origin-destination level.

4.1 Selection Bias

A potential critique of our empirical approach is 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, if anything, will 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 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

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.

logarithm of the number of households moving across an MSA pair in a given year on year and 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 Indeed, our results in section 5.4 will show that households who move across highly covarying cities are more likely to be homeowners. 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. For simplicity, consider a version of the empirical model in Equation 7 that relates house value (y_h) with a single covariate, cross-market house price covariance, which is denoted as above by c_{ij} . H_h is an indicator variable for whether household h owns its home.

$$y_h = \beta_0 + \beta_1 c_{ij} + \eta_h$$

$$H_h = 1 [\gamma_0 + \gamma_1 c_{ij} + \phi_h > 0]$$

$$\text{corr}(\eta_h, \phi_h) > 0$$

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

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

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.

the covariance rises. This biases estimates of both the conditional mean and conditional variance.

If y_h 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 y_h and ϕ_h are conditionally bivariate normal, although it does not hold for all distributions in general. The distribution of home values in our sample, conditional on covariates, is approximately normal. Consequently, under a conditional normality assumption on ϵ_h , any selection biases our estimates toward zero.

5 Results

5.1 Conditional Variance Estimates

We present our results in three parts. The first section relates the conditional variance of house values to the covariance between origin and destination MSA pairs. The second section looks at the effect of covariance on conditional mean housing demand, while the third 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 the covariates so the estimated coefficients can be interpreted as the marginal effects of a one-standard deviation change in the covariate. 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. 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 first period 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 col-

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.

umn (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.

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” 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 second-period housing spending, along with a 95 percent confidence interval, is shown in Figure 2. 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. 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 results, equations 4 and 5, prescribe a linear effect only through the use of a Taylor approximation.

5.2 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

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.

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 4, 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 4 literally and interact covariance with factors that shift q_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.²⁸ 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 ($q_{2p}(\cdot)$ and $q_{2w}(\cdot)$) in Equation 4. 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 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_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,

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.

²⁸This regression is very similar to the conditional mean regressions discussed below, although we exclude covariance from the set of covariates. Since we include destination fixed effects, using log house value rather than house value itself guarantees that predicted differences are differences in quantity rather than just price, under the maintained assumption that house prices per unit quantity are constant within an MSA in a given year.

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 (q_1) substantially and statistically significantly amplifies the effect of covariance, by 4.1 percentage points. These results match our predictions based on Equation 4.

Since both income and predicted q_1 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_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.

5.3 Conditional Mean Estimates

Our theoretical derivation indicates that the direct effect of covariance on mean housing spending should be similar to the effect on the variance, which is precisely what we find. The specifications reported in this section parallel the tables from the conditional variance section and are similar to the first stage of the conditional variance estimation procedure, with the sole difference being that we use log housing values as the dependent variable. As with the conditional variance estimation routine, using logs allows us to separate quantity from price, since the destination fixed effects will absorb the effects of differences in prices as long as they are additively separable from quantity. Table 5 reports the results of the conditional mean regression with a varying set of controls. Column (1) shows the coefficients from a regression of house values on covariance and origin and destination fixed effects, while column (2) includes our standard set of covariates. Column (3) adds controls for current income and column (4) adds our proxy for *ex ante* expected covariance.

In all cases, higher covariance between origin and destination leads to lower mean housing expenditures in the destination city, even after including

the full set of covariates and origin and destination MSA fixed effects. A one standard deviation higher expected covariance yields anywhere from a 2.6 to 0.9 percent reduction in housing spending in the destination, depending on the specification. All of these coefficients are fairly tightly estimated and are statistically distinguishable from zero.

According to equation 5, the impact of covariance on mean house spending should be concentrated amongst households who own more housing before moving. In Table 6, we interact covariance with income or predicted house size (q_1). With separate origin and destination fixed effects and a full set of household controls, including expected covariance, we find large and statistically significant effects of the ratio of income to origin price (1) and predicted q_1 (2) on the impact of covariance on mean housing expenditures. A one standard deviation increase in the income-price ratio more than doubles the effect of covariance, from -0.011 to -0.024, while a one standard deviation increase in predicted housing has a similar effect. These interactions are nearly identical when origin x destination fixed effects are included in columns (3) and (4).

We also allow the effect of covariance on mean destination housing consumption to be nonlinear by estimating a generalized additive model. As with the regressions, we put log house value on the left hand side, but the model is otherwise identical to the first-stage GAM used in the conditional variance estimates. In contrast with Figure 2, Figure 3 shows that the effects above and below the mean are roughly the same — both around -0.011 — although the negative slope is more precisely estimated above the mean. Overall, the pattern of results are quite similar for the conditional mean and conditional variance regressions, as our theory predicts. They are separably estimable because the variance regressions already control for mean housing spending, so the conditional variance results are not due to the differences in conditional mean.

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 dependent variable, and the same sets of covariates as in the conditional variance estimates.²⁹ The

²⁹We 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

estimation sample contains approximately twice as many households as our sample of homeowners.

The baseline results are reported in Table 7. 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 homeownership rate is 50 percent, so 3 percentage points corresponds to a 6 percent increase.³⁰ 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 8 repeats the strategy from Table 4 of interacting covariance with factors that shift first-period housing quantity q_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_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 cities at the high end of the covariance distribution. The first column of Table 9 uses our estimates to calculate the percentage effect of a one standard deviation increase in covariance on post-move housing consumption variance.

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).

³⁰The homeownership rate in the selected sample of inter-MSA migrants is substantially lower than the national homeownership rate because renters are much more likely to move than homeowners.

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 4.³¹ 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 1, 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 just 40 percent the variance in post-move consumption relative to 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.

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 con-

³¹For ease of explanation, we do not use the nonlinear estimates from the GAM described above and shown in Figure 2. The patterns described here would be even more stark if we did.

sumption 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 non-linearity and heterogeneity across groups. Allowing the estimated coefficient to vary nonlinearly 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.

We find additional evidence consistent with the model. On average, households who face higher covariances tend to spend less on housing after a move because of the convexity of the housing demand function. They also

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.

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.
- 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.
- 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. and Joseph Gyourko**, “Housing Dynamics,” 2007. mimeo.
- 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**, “The Effects of Price Uncertainty on Housing Demand: Empirical Evidence from U.S. Markets,” 2008. mimeo.
- , “Hedging House Price Risk in the Presence of Lumpy Transaction Costs,” *Journal of Urban Economics*, 2008, *64* (2), 270–287.
- Hastie, T.J. and R.J. Tibshirani**, *Generalized Additive Models*, Chapman and Hall/CRC, 1990.
- 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

Case	State of the world		Elasticity parameters	
	If house prices in city B changed by:	...house prices in city A must have changed by:	Cobb-Douglas $\epsilon_w = 1$ $\epsilon_p = -1$	Rosen (1979) $\epsilon_w = 0.75$ $\epsilon_p = 1$
Perfect covariance	+20% -20%	+20% -20%	0% 0%	-5% +5%
Perfect negative covariance	+20% -20%	-20% +20%	-40% +40%	-35% +35%

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	Estimation Sample 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			0.6604	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 Percentage Effect of Covariance on Housing Expenditure Mean

Variable	(1)	(2)	(3)	(4)
Covariance (standardized)	-0.026 (0.005)	-0.017 (0.004)	-0.009 (0.003)	-0.011 (0.003)
Household income (\$1000's)			0.006 (0.000)	0.006 (0.000)
Household income -squared			-0.000007 (0.000000)	-0.000007 (0.000000)
Exp. Covariance (standardized)				0.005 (0.004)
Observations	100851	100851	100851	100851
R^2	0.217	0.398	0.509	0.509
Origin FE	X	X	X	X
Destination FE	X	X	X	X
Household controls		X	X	X

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 Percentage Effect of Covariance on Housing Expenditure Mean

Variable	(1)	(2)	(3)	(4)
Covariance (standardized)	-0.011 (0.003)	-0.011 (0.003)		
Covariance x Income / Origin Price	-0.013 (0.003)		-0.014 (0.004)	
Covariance x Predicted House Size		-0.018 (0.003)		-0.017 (0.003)
Observations	100851	100851	100851	100851
R^2	0.509	0.509	0.577	0.577
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: 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 8: 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 9: 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.

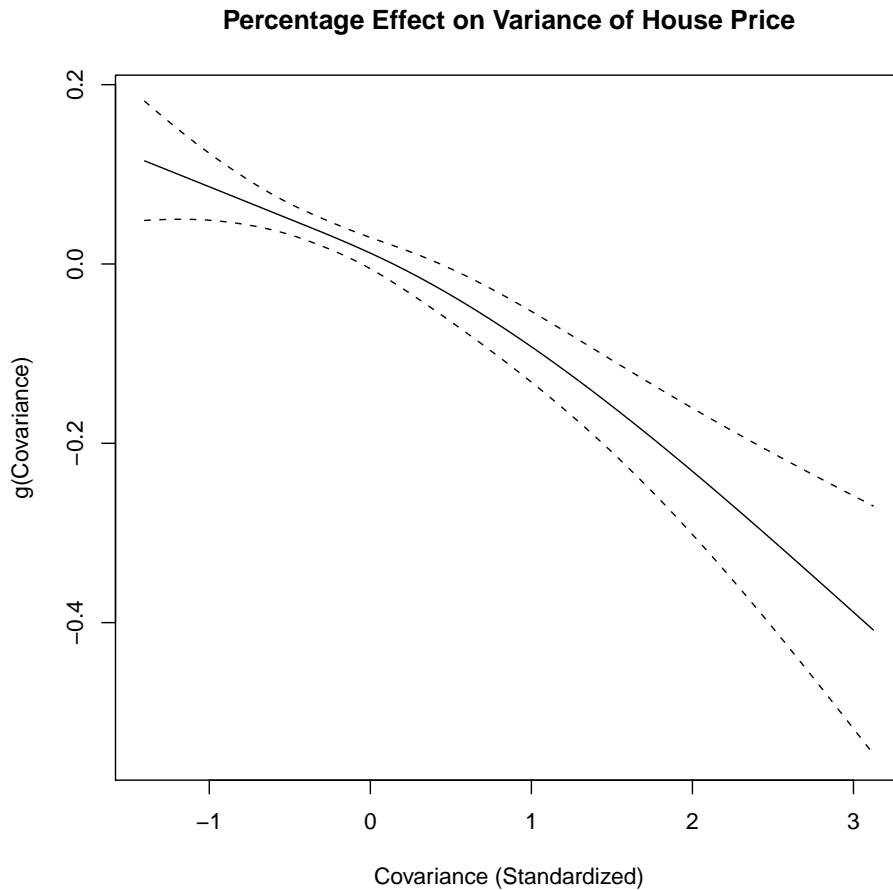
Figure 1:

Histogram of 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 2:



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.

Figure 3:

Percentage Effect on Mean 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.