

OWNER-OCCUPIED HOUSING AS A HEDGE AGAINST RENT RISK*

TODD SINAI AND NICHOLAS S. SOULELES

The conventional wisdom that homeownership is very risky ignores the fact that the alternative, renting, is also risky. Owning a house provides a hedge against fluctuations in housing costs, but in turn introduces asset price risk. In a simple model of tenure choice with endogenous house prices, we show that the *net* risk of owning declines with a household's expected horizon in its house and with the correlation in housing costs in future locations. Empirically, we find that both house prices, relative to rents, and the probability of homeownership increase with net rent risk.

Economists typically treat houses like standard financial assets. This leads to the conventional view that owning a house is quite risky. Since house prices are volatile and homeowners on average allocate 27 percent of their net worth to their house [Poterba and Samwick 1997], fluctuations in house prices can have a sizable effect on homeowners' balance sheets. Further, prior studies have emphasized that changes in housing wealth can lead to significant changes in homeowners' consumption (e.g., Case, Quigley, and Shiller [2003]), and overinvestment in housing can distort their financial portfolio allocations (e.g., Brueckner [1997], Flavin and Yamashita [1998], Fratantoni [1997], and Goetzmann [1993]).

This focus on the asset price risk of home owning neglects the fact that all households are in effect born "short" housing services, since they have to live somewhere. Households that do not own must rent, purchasing their housing services on a spot market, and thus subjecting themselves to annual fluctuations in

* We are grateful to Andrew Abel, Edward Glaeser, Joao Gomes, Joseph Gyourko, Matthew Kahn, Christopher Mayer, Robert Shiller, Jonathan Skinner, Amir Yaron, Bilge Yilmaz, three anonymous referees, and participants in seminars at the AEA/AREUEA 2001 annual meetings, NBER, Syracuse University, University of Wisconsin, University of British Columbia, University of California, Berkeley, the Federal Reserve Bank of Philadelphia, and The Wharton School for their helpful comments and suggestions. James Knight-Dominick and Daniel Simundza provided excellent research assistance. Sinai acknowledges financial support from the Research Scholars Program of the Zell/Lurie Real Estate Center at Wharton. Souleles acknowledges financial support from the Rodney L. White Center for Financial Research and a NYSE Fellowship through the Rodney L. White Center. Address correspondence to Todd Sinai, The Wharton School, University of Pennsylvania, 1465 Steinberg Hall-Dietrich Hall, 3620 Locust Walk, Philadelphia, PA 19104-6302. Phone: (215) 898-5390. E-mail: sinai@wharton.upenn.edu.

rent. Owners, by contrast, avoid this rent uncertainty by buying a long-lived asset that delivers a guaranteed stream of housing services for a known up-front price. Unlike standard assets, houses effectively pay out annual dividends equal to the ex post spot rent, and so provide a hedge against rent risk. Hence considering asset price risk in isolation fails to account for households' entire risk position. For example, if rents rise, all households would experience an increase in their implicit future rent "liability." But homeowners would enjoy a compensating increase in the price of their house, since its dividend payments would increase commensurately. Such offsetting effects reduce the overall wealth effects from changes in house prices.

Homeowners still face asset price risk, but only when they move (or die) and sell their house. Since that risk comes in the future, it can be small in present value. The key reason that there is any price risk at all is that houses "outlive" their owners' residence spells—the hedge is in place too long—and the houses' value at the end of the spells is risky. If residence spells were infinite (or in a dynastic setting, if descendants live in the same houses as their parents), homeownership would not be risky at all, since there would be no sale price risk.

We show, in a simple tenure choice model with endogenous house prices, how the demand for owning trades off the rent and asset price risks. Which risk dominates on net is largely determined by households' expected length of stay (horizon) in their houses and whether they move to correlated housing markets. First, for households with longer horizons, the rent risk is more likely to dominate, increasing the demand for home owning, since these households experience a greater number of rent fluctuations and the asset price risk arrives further in the future and so is more heavily discounted. Further, the *net* rent risk increases in magnitude with the interaction of rent volatility and horizon, so the demand for owning increases faster with rent volatility for households with longer horizons. Second, greater spatial correlation in house prices across housing markets reduces the total magnitude of asset price risk because, at the time of a move, the sale and repurchase prices are more likely to offset, extending households' effective horizons. Also, insofar as house prices are persistent over time, the purchase price of a future house is partially hedged by its own subsequent sale price, which also reduces the total price risk. The conventional view that homeownership is very risky can be seen as a special case of this model,

under the assumption that households' horizons are short (ignoring rent risk altogether) and housing markets are uncorrelated.¹

We find empirical evidence consistent with the model's predictions that the demand for homeownership should increase with local rent volatility, households' expected horizon, and their interaction. Depending on the elasticity of supply of owned housing units, the hedging demand for owning may show up in a higher homeownership rate, higher house prices, or both. Using household-level data, we find that the probability of homeownership increases faster with rent volatility for long-horizon households than for short-horizon households. The difference between the probability of homeownership for a household with an above-the-median expected horizon and that of a below-the-median household is up to 5.4 percentage points greater in metropolitan statistical areas [MSAs] with high rent variance than in MSAs with low rent variance. Further, the sensitivity to rent risk is greatest for households that face bigger housing gambles, in that typical rents in their MSA comprise a relatively large portion of their annual income.

We also find that house prices, normalized by rents, capitalize a premium for avoiding net rent risk. At the MSA level, a one standard deviation increase in rent variance, holding expected rents constant, raises house prices by 2 to 4 percent. The price-to-rent ratio also increases with expected future rents, just as a price-earnings ratio for stocks should increase with expected future earnings.

The remainder of this paper proceeds as follows. In Section I we present a stylized model of tenure choice in the presence of both rent risk and house price risk. Section II describes our data sources and variable construction. The empirical methodology

1. The traditional user cost literature, e.g., Rosen [1979], Hendershott and Slemrod [1983], and Poterba [1984], does not consider risk at all. More recent work neglects the trade-off between rent and price risk. Skinner [1989] and Summers [1981] consider only price risk, and Ben-Shahar [1998] only rent risk. Rosen, Rosen, and Holtz-Eakin [1984] incorporate both risks but do not allow for an endogenous relation between house prices and rent. The recent portfolio choice literature treats housing like other financial assets and models its price risk, along with other sources of household risk, but neglects the tenure decision and rent risk. Davidoff [2003] measures asset price risk by how much house prices covary with labor income. See also, e.g., Cocco [2000], Haurin [1991], Campbell and Cocco [2003], Flavin and Nakagawa [2003], and Berkovec and Fullerton [1992]. In work subsequent to this paper, Ortalo-Magne and Rady [2002a] develop an extended version of our framework that examines the implications of the covariance between rents and earnings. Henderson and Ioannides [1983] focus on incentive compatibility between landlords and tenants. Chan [2001], Genesove and Mayer [1997], and Stein [1995] consider the effects of house price declines on mobility.

and results are reported in Section III. Section IV briefly concludes.

I. A MODEL OF TENURE CHOICE AND HOUSE PRICES

This section presents a stylized model of tenure choice in which the cost of securing housing services is uncertain and house prices are endogenous. We show how equilibrium house prices, which will fully capitalize the demand for homeownership, reflect the present value of expected future rents plus a risk premium that compares the rent risk avoided by owning with the asset price risk incurred. We then demonstrate how the trade-off between the rent and asset price risks varies with households' horizon and the volatility and spatial correlation of housing markets.

To emphasize the key aspects of this trade-off, we make a number of simplifying assumptions. We consider *ex ante* identical, risk-averse households that live in two locations, first *A* then *B*, for *N* years each, after which they die.² At birth, labeled year 0, a household chooses whether to be a homeowner in both locations or to rent in both locations. Its desired quantity of housing services is normalized to be one unit and is the same in each location. For convenience, rental units and owner-occupied units, in fixed supply and together equal to the number of households, both provide one unit of housing services.³ When the households leave *A* or *B*, they are replaced by a similar generation of households.

The cost of housing services, represented by the spot rent, fluctuates due to exogenous shocks to the underlying economy and housing market.⁴ We allow for correlation in rents across the

2. The occurrence and timing of the move are assumed to be exogenous and known in advance. If transactions costs are greater for owners, uncertainty about the total number of moves could generate an additional risk of owning. Allowing for endogenous moves, as in Ortalo-Magné and Rady [2002b], could reduce the risks of both renting and owning, as households could move to a location with lower housing costs.

3. The results below can be generalized to allow the services from an owner-occupied house to exceed those from renting, perhaps due to agency problems. In practice, rent risk might also reduce the desired size of housing space (the intensive margin). That effect would reinforce the rent versus own (extensive) margin that we analyze empirically.

4. Any number of local economic conditions fluctuate over time and across space, from the success of locally concentrated industries to increased immigration. In an earlier version of this paper, we found that rent volatility in a market varies with, among other things, volatility in the local unemployment rate interacted with the inelasticity of supply of housing (proxied by regulatory constraints on building).

two locations, as well as persistence over time, by assuming that the spot rents in the two locations follow correlated AR(1) processes: $r_n^A = \mu^A + \varphi r_{n-1}^A + k(\eta_n^A + \rho \eta_n^B)$ and $r_n^B = \mu^B + \varphi r_{n-1}^B + k(\rho \eta_n^A + \eta_n^B)$, where $\varphi \in [0,1]$ measures the persistence of rents, μ^A and μ^B measure the expected level or growth rate of rents (depending on φ), and the shocks η^A and η^B are independently distributed $\text{IID}(0, \sigma_A^2)$ and $\text{IID}(0, \sigma_B^2)$.⁵ ρ parameterizes the spatial correlation in rents (and, endogenously, in house prices) across the two locations, with $\rho = 0$ implying independence and $\rho = 1$ implying perfect correlation. To control the total magnitude of housing shocks incurred as ρ varies, the scaling constant κ can be set to $1/(1 + \rho^2)^{1/2}$.

Assuming perfect capital markets and known, exogenous lifetime wealth, W , the representative household will choose its tenure mode to maximize its expected utility of wealth net of total housing costs, or equivalently to minimize its total risk-adjusted housing costs. The ex post cost of renting, discounted back to year 0, is simply the present value of the rents faced in A and then B , assuming the move takes place in year N and leases fix rents for one year:⁶

$$C_R \equiv r_0^A + \sum_{n=1}^{N-1} \delta^n \tilde{r}_n^A + \sum_{n=N}^{2N-1} \delta^n \tilde{r}_n^B,$$

where δ is the discount factor. The initial rent r_0^A is observed at time 0, but the future rents, identified with tildes, are stochastic.

For owners, the ex post cost of housing is the difference between the discounted purchase and sale prices of both houses: $C_O \equiv P_0^A + \delta^N(\tilde{P}_N^B - \tilde{P}_N^A) - \delta^{2N}\tilde{P}_{2N}^B$. Only the initial purchase price in A , P_0^A , is known as of year 0.⁷ We will show below that

5. We take the spot rent processes as given. Whatever their ultimate determinants, the model correctly specifies the endogenous relationship that results between rents and house prices. This approach is analogous to other asset-pricing models. For instance, in term structure models of long versus short maturity bonds, the process for short rates (analogous to our rental rates) is the exogenous input into the model. In models of stock prices, the input is the process for a firm's cash flows, and the stochastic price of its stock at sale is analogous to our house sale price.

6. Genesove [1999] reports that 97.7 percent of all residential leases are for terms of one year or less. Also, one cannot purchase a "rent swap" to exchange variable rents for fixed rents. We suspect that such long-term and swap contracts do not exist because of agency problems. For example, renters could unilaterally exit a swap by moving if rents fell.

7. For simplicity, we abstract from other factors that affect homeownership and rental costs, such as the tax treatment of homeownership, maintenance, and depreciation. (See, e.g., Hendershott and Slemrod [1983] and Gyourko and Tracy [2004].) Such factors may affect the relative cost of owning and renting, but they would not qualitatively change the comparative statics at issue here regarding the effects of cross-sectional variation in horizon and rent volatility.

house prices can be expressed as linear functions of contemporaneous rents, and so the future house prices \tilde{P}_N^A , \tilde{P}_N^B , and \tilde{P}_{2N}^B endogenously fluctuate with the respective rent shocks.

Since the supply of housing is fixed, the demand for owning will be entirely capitalized into house prices. We focus on prices in *A*, which will incorporate the future risks from moving, and assume that in equilibrium these prices adjust to leave each generation of households ex ante indifferent at the start of their lives between renting and owning. Hence P_0^A sets $E_0U(W - C_R) = E_0U(W - C_O)$. With a stationary environment there is no loss of generality in expressing the results just for year 0.

The resulting equilibrium price consists of the expected present value of future rents, *PV*, plus the total risk premium households are willing to pay to own rather than rent, *RP*.⁸

$$(1) \quad \begin{aligned} P_0^A &= PV(r_0^A) + RP(N, \sigma_A^2, \sigma_B^2, \rho, \varphi, \alpha) \\ &= PV(r_0^A) + \frac{\pi_R - \pi_O}{1 - \delta^N} - \frac{\delta^N(\pi_R^B - \pi_O^B)}{1 - \delta^N}. \end{aligned}$$

The *PV* term arises because houses provide a constant housing service flow in perpetuity, of value equal to the spot rents saved. Just as a price-earnings ratio increases with expected earnings growth, the difference between the house price P_0^A and the current rent r_0^A increases with expected future rents in *A*, and so with the trend μ^A , as well as the persistence of rents, φ :

$$(2) \quad PV(r_0^A) \equiv \frac{1}{1 - \delta\varphi} \left(r_0^A + \mu^A \frac{\delta}{1 - \delta} \right).$$

The total risk premium *RP* compares the risks of renting and owning. The risk premium for renting, π_R , measures the risk associated with the cost C_R of renting. Because owning provides the benefit of avoiding the rent shocks, in equilibrium the rent risk premium is bid into house prices, so π_R enters equation (1) with a positive sign. This has the important implication that rent risk tends to increase the demand for home owning, ceteris pari-

8. In short, equation (1) comes from equating the certainty-equivalent utilities of renting and owning, $U(W - E_0C_R(r_0^A, r_0^B, \mu^A, \mu^B) - \pi_R(\eta_1^A, \dots, \eta_{2N-1}^A, \eta_1^B, \dots, \eta_{2N-1}^B)) = U(W - E_0C_O(P_0^A, r_0^A, r_0^B, \mu^A, \mu^B) - \pi_O(\eta_1^A, \dots, \eta_{2N}^A, \eta_1^B, \dots, \eta_{2N}^B))$ and solving out for P_0^A , after recursively expressing each rent in location *j* as $r_t^j = \varphi^t r_0^j + \mu^j \sum_{i=1}^t \varphi^{i-1} + \sum_{i=1}^t \varphi^{t-i} k(\eta_i^j + \rho \eta_i^{-j})$, and future prices P_t^j as a function of r_t^j . Prices in *B* are assumed to leave households indifferent between owning and renting in *B*, once they get to *B*. Equivalently, they can be determined by another sequence of households with synchronized horizons of *N* who live only in *B*.

bus. The risk premium for owning, π_O , measures the risk associated with the costs C_O of owning. Owners incur asset price risk, which reduces the demand for owner-occupied housing, *ceteris paribus*. Hence π_O enters equation (1) with a negative sign.

Overall, P_0^A depends on $\Delta\pi \equiv (\pi_R - \pi_O)$, the difference between the risk premiums for renting and owning (in *A* and *B*). If renting is riskier on balance than owning, then the sign of the *net* risk premium $\Delta\pi$ would be positive (and vice versa). In this case, unless spot rents are riskless ($\sigma_A^2 = \sigma_B^2 = 0$) or households are risk neutral ($\alpha = 0$), as in Poterba [1984], households would bid up the house price P_0^A relative to the *PV* term because of the hedging benefit that the house provides against rent risk. P_0^A also capitalizes $\Delta\pi^B \equiv (\pi_R^B - \pi_O^B)$, a net risk premium for renting versus owning in *B* that in equilibrium P_0^A inherits from house prices in *B*.

We present analytic solutions for each component of *RP* in equation (1), and then proceed directly to how they vary with the factors highlighted in the introduction, especially the horizon *N*, local rent volatility σ_A^2 , and spatial correlation ρ . To highlight the underlying intuition, we initially present the simplest case in which rents are IID and spatially uncorrelated, with $\varphi = 0$ and $\rho = 0$, and subsequently introduce the effects of spatial correlation, allowing $\rho > 0$. The general case, allowing for persistent rents with $\varphi > 0$, is left to the Appendix.

The rent risk premium π_R measures the total risk arising from the sequence of annual rent shocks. In the special case with $\varphi = 0$ and $\rho = 0$, it can be approximated by

$$(3) \quad \pi_R \approx \frac{\alpha}{2} \left(\sigma_A^2 \sum_{n=1}^{N-1} (\delta^n)^2 + \sigma_B^2 \sum_{n=N}^{2N-1} (\delta^n)^2 \right),$$

where α is household risk aversion. The first summation reflects the “early” rent shocks η_1^A to η_{N-1}^A incurred while living in *A*, with the shocks coming later discounted more heavily. The second summation reflects the “late” rent shocks η_N^B to η_{2N-1}^B incurred in *B*. More generally, if rent shocks are persistent ($\varphi > 0$) or spatially correlated ($\rho > 0$), then π_R further increases with additional terms reflecting the resulting extra rent risk, as shown in Appendix equation (A1).

The risk premium for owning π_O measures the total asset price risk resulting from the three future housing transactions: the sale of the house in *A*, the simultaneous purchase of a house in *B*, and the later sale of the house in *B*. When $\varphi = 0$ and $\rho = 0$,

$$(4) \quad \pi_O \approx (\alpha/2)(\delta^{2N}(\sigma_A^2 + \sigma_B^2) + \delta^{4N}\sigma_B^2).$$

In this case the price risks come entirely from the impact of the *contemporaneous*, local rent shocks on the three future house prices: η_N^A on P_N^A , η_N^B on P_N^B , and η_{2N}^B on P_{2N}^B . With spatial independence, the risk from the buy-in-*A* and sell-in-*B* transactions is the sum of the risks of the individual transactions ($\sigma_A^2 + \sigma_B^2$), discounted by δ^{2N} since they take place in year N . The terminal sale risk in *B* is discounted by δ^{4N} since it arrives another N years later. More generally, if rents are persistent or spatially correlated, the future prices will also fluctuate with the *preceding* local rent shocks (if $\varphi > 0$) and with the corresponding shocks from the other location (if $\rho > 0$), as shown in equation (A2).

Finally, in equilibrium P_0^A also indirectly capitalizes risk premiums that are embedded in the future house prices. In equation (1) the factor $1/(1 - \delta^N)$ reflects the fact that both the purchase price P_0^A and the sale price P_N^A include the same total risk premium RP , since both current and *future* owners must be indifferent between owning and renting. When the initial owners pay P_0^A , they take into account that they will recoup this premium, albeit at a discount, N periods later on selling for P_N^A . Similarly, the term $\Delta\pi^B = (\pi_R^B - \pi_O^B)$ takes into account the fact that the future purchase price in *B*, P_N^B , itself capitalizes a premium for the net rent risk avoided by owning while in *B*. That premium in P_N^B is in equilibrium inherited by P_0^A . For instance, if $\Delta\pi_B$ increases, increasing P_N^B , to compensate P_0^A must decrease sufficiently to keep the initial owner indifferent between renting and owning. Fixed moving/transactions costs would have an analogous effect on P_0^A . *Ceteris paribus*, P_0^A would fall by the present value of the costs in order to compensate homeowners. When $\varphi = 0$ and $\rho = 0$, the net risk premium $\Delta\pi^B$ in *B* is simply the present value of the $N - 1$ rent risks (η_N^B to η_{2N-1}^B) avoided in *B*, less the discounted sale price risk from P_{2N}^B . When capitalized in P_0^A , this premium is additionally discounted by δ^N since P_N^B is not paid until year N :

$$(5) \quad \delta^N(\pi_R^B - \pi_O^B) \approx \delta^N \frac{\alpha}{2} \sigma_B^2 \left(\sum_{n=1}^{N-1} (\delta^n)^2 - (\delta^N)^2 \right).$$

We now turn to how the demand for owning, as reflected in P_0^A , varies with the key parameters. Using the general versions of the risk premiums from the Appendix, we show that as the horizon N or spatial correlation ρ increases, the demand for

homeownership typically increases. The average effect of local rent volatility σ_A^2 is ambiguous—it depends on the horizon of the marginal household—but demand will nonetheless generally increase with the interaction of the rent volatility and horizon ($N \times \sigma_A^2$).

First, the demand for homeownership generally increases with a household's expected horizon in its home, so $\partial P_0^A / \partial N > 0$. This result is driven by how the net risk premium, $\pi_R - \pi_O$, varies with N . The π_R component increases with N in equations (3) and (A1), in the symmetric case with $\sigma_B^2 = \sigma_A^2$. This is because, as the horizon lengthens, the renter faces more rent shocks, increasing the total amount of rent risk. π_O tends to decline with N in equations (4) and (A2), because with a longer horizon the price risks come later in time and are discounted more heavily. In the limit, with infinite residence spells, π_O converges to zero while π_R is positive. In this case owning would be completely riskless, since a household would buy at the initial, known market price (P_0^A) and never face any price or rent risk. The house would then provide a perfect hedge against fluctuations in housing costs. Since utility is determined by the housing service flow rather than by the house price, then absent borrowing and collateral constraints, the unrealized house price fluctuations would not have any effect on household utility.⁹ By contrast, with finite horizons the future sale and new purchase of houses leads to asset price risk ex post, since the houses outlast their owners' residence spells.

As a result, the net risk premium $\Delta\pi$ tends to increase with N , as the rent risk becomes increasingly large relative to the price risk. While the other premiums that are capitalized into P_0^A indirectly through the future sale and purchase prices can work in the opposite direction, for reasonable degrees of persistence and spatial correlation, P_0^A generally increases with N . In fact, for completely persistent rent shocks, $\varphi = 1$, P_0^A is monotonically increasing in N . In our MSA-level data, rents exhibit substantial persistence, with φ about 0.85 (see also Case and Shiller [1989]), and a high degree of spatial correlation, with ρ averaging about 0.50. Using these values, along with $\sigma_B^2 = \sigma_A^2$ and a discount factor of $\delta = 0.96$, as our benchmark parameters, P_0^A rises steeply with N for up to thirteen years, then slightly declines and plateaus.

9. When a house is a mechanism for borrowing, declines in house prices could potentially reduce a household's consumption and welfare even if its horizon is infinite [Zeldes 1989; Engelhardt 1996; Jappelli, Pischke, and Souleles 1998; Gross and Souleles 2002].

Second, the effect of local rent volatility on demand, $\partial P_0^A / \partial \sigma_A^2$, also generally increases with horizon, with $\partial^2 P_0^A / \partial \sigma_A^2 \partial N > 0$. That is, the demand for owning typically increases with the interaction of local rent variance and horizon. To see why, note that σ_A^2 and σ_B^2 enter all the risk premiums in RP in equation (1) multiplicatively and separately. Hence, $\partial P_0^A / \partial \sigma_A^2$ is proportional to the half of RP that is multiplied by σ_A^2 , and so the effect of the horizon N on $\partial P_0^A / \partial \sigma_A^2$ is parallel to the effect discussed above of N on RP and P_0^A . For our benchmark parameters, $\partial P_0^A / \partial \sigma_A^2$ rises steeply with N for 25 years, and again then slightly declines and plateaus. Therefore, $\partial^2 P_0^A / \partial \sigma_A^2 \partial N$ is positive for most relevant horizons. Our empirical work will emphasize this interaction result because it will be the most convincingly identified in our data.

Because $\partial P_0^A / \partial \sigma_A^2$ generally rises with N , greater rent volatility σ_A^2 tends to increase demand when N is large, as renting at some point becomes riskier than owning. But for small N , $\partial P_0^A / \partial \sigma_A^2$ can be negative. For the benchmark parameters, in our stylized model $\partial P_0^A / \partial \sigma_A^2$ is positive as long as the horizon N is above four years. More generally, the *average* effect of rent volatility on demand might possibly be small in magnitude and even ambiguous in sign, depending on the horizon of the marginal household.¹⁰

Finally, greater spatial correlation ρ in house prices reduces the net risk of owning, justifying higher house prices, *ceteris paribus*. In equations (4) and (A2), with $\sigma_B^2 = \sigma_A^2$, π_O monotonically declines with increases in ρ , and as a result $\partial P_0^A / \partial \rho$ is positive. This occurs because, to the degree that housing costs in A and B are correlated, the price risk from selling the house in A is undone by the price paid for the new house in B . To see this, equation (4) can be generalized to allow for spatial correlation in housing costs, with $\rho > 0$:

$$(6) \quad \pi_O \approx \frac{\alpha}{2} \kappa^2 [\delta^{2N} (1 - \rho)^2 (\sigma_A^2 + \sigma_B^2) + \delta^{4N} (\rho^2 \sigma_A^2 + \sigma_B^2)].$$

The resulting cross-sectional “natural” hedge is captured by the new $(1 - \rho)^2$ factor, which reflects the offsetting sale and purchase risks at N . With perfect correlation of $\rho = 1$, the shocks to P_N^A and P_N^B completely cancel each other out, because the sell-in- A /buy-

10. The conclusions are qualitatively similar on setting $\sigma_A^2 = \sigma_B^2 = \sigma^2$ and analyzing the effects of “national” increases in rent volatility on demand: $\partial P_0^A / \partial \sigma^2$ and $\partial^2 P_0^A / \partial \sigma^2 \partial N$.

in-*B* transaction at N is a complete wash. For example, with $\rho = 1$ and $\sigma_A^2 = \sigma_B^2$, equation (6) reduces to $\pi_O = (\alpha/2)(\delta^{4N}\sigma_B^2)$, reflecting only the subsequent terminal sale price risk from P_{2N}^B . (The new $\rho^2\sigma_A^2$ term in equation (6) reflects the contribution of the additional location-*A* rent shock η_{2N}^A on P_{2N}^B , which is offset by the scaling factor $\kappa = 1/(1 + \rho^2)^{1/2}$ when $\sigma_B^2 = \sigma_A^2$.)

More generally, as ρ increases, the sale/buy transaction becomes more of a wash sale, which effectively extends the owner's horizon, and so reduces the asset price risk. Thus, even with finite residence spells, if an infinitely lived household moves across perfectly correlated locations, or in a dynastic setting, if each generation passes on the house for its descendants to live in, the effective horizon would again be infinite. Even if a household sells its house at death, if it bequeaths the proceeds to its descendants and they use the inheritance to buy another house in the same or correlated market, the effective horizon is again longer.

Allowing rents to be persistent with $\varphi > 0$ introduces another, intertemporal natural hedge, evident in equation (A2). For example, if prices in *B* turn out to be high when the household buys in year N , the offsetting sale price in year $2N$ would also be expected to be high, dampening the overall effect on the cost of owning.

The conventional view that housing is very risky can be seen as a special case of this model, under the assumption of short horizons N (or ignoring rent risk altogether) and low spatial correlation ρ . In this case, the asset price risk is very salient, and the rent risk is relatively less important. But empirically, spatial correlation is significant, and many households have long horizons, in which case the asset price risk becomes smaller and the rent risk more important. For households with even moderate horizons, the rent risk may dominate.

This analysis assumed that households do not adjust the size of their housing consumption bundles on moving. If we allowed such adjustments, house price fluctuations could actually increase owners' utility. Consider, in a partial equilibrium setting without adjustment costs, the consumption bundles that a household could choose if house prices fluctuate, using a revealed preference argument. If house prices rise, an infinitely lived household could stay in its existing house and maintain its original level of utility; or, it could sell its appreciated house, purchase a smaller house, and use the remaining proceeds to consume more nonhousing goods, if doing so makes it better off. Conversely, if house prices fall, the household could stay in the

existing house and maintain its original level of utility, or substitute toward housing by consuming fewer nonhousing goods and buying a bigger house. In either case, by revealed preference, the household is no worse off and might be better off.¹¹

The analysis also suggests that the aggregate wealth effect from house price fluctuations is likely to be relatively small, in the absence of important borrowing or collateral constraints. Equation (1) implies that, absent changes in risk premiums and discount rates, increases in house prices reflect a commensurate increase in the present value of expected future rents, which increases the cost of fulfilling households' short position in housing services. For homeowners with infinite horizons, this increase in implicit liabilities would exactly offset the increase in the house value (their long position), leaving their effective expected net worth unchanged. Even for homeowners with finite horizons who move to uncorrelated markets, every housing transaction is just a transfer between a buyer and a seller. If the propensity to consume out of wealth is similar on average across buyers and sellers, then any resulting wealth effects from house price fluctuations would tend to wash out. This might help explain why studies of the propensity to consume out of housing wealth find smaller effects at the aggregate level than at the micro level.¹²

II. DATA AND VARIABLE CONSTRUCTION

To test the implications of the model, we combine data on rents, house prices, and homeownership from multiple sources. Starting at the MSA level, average annual rents come from surveys of "Class A" (top-quality) apartments by Reis, a commercial real estate information company. We derive annual house prices by MSA by inflating each MSA's median house price from the 1990 Census by the corresponding growth rate of the Freddie Mac repeat-sales house price index. Both rents and prices are converted to real dollars using the CPI excluding shelter. After

11. We thank Edward Glaeser for this example.

12. See Case, Quigley, and Shiller [2003] and Skinner [1996]. Indeed, bringing renters back into the picture can potentially reverse the conventional logic regarding wealth effects. Consider an increase in house prices that is due to an exogenous increase in expected future rents. Renters (either current renters or future renters depending on the timing of the rent increases) would experience a negative wealth effect due to the increased housing costs. So it is possible for aggregate consumption to *decline* at the same time that house prices rise, especially if the asset-price effect on the buying and selling households is approximately a wash. See also concurrent research by Bajari, Benkard, and Krainer [2003].

combining the rent and price data, we have complete observations on 44 MSAs from 1981 through 1998.

We use these data to measure the expected growth and volatility of both rents and prices by MSA. For an MSA in a given year t , the expected growth rate is the average change in log real rents, or log real house prices, over the prior nine-year period, $t - 1$ to $t - 9$. These lags shift the effective sample period to $t = 1990$ – 1999 . Rent variance is computed using the within-MSA annual differences between the actual log rent and the calculated average growth rate over the prior nine years, and is expressed as a percentage of the base rent. Using rent volatility to measure rent risk is analogous to using income volatility to measure income risk, as in the portfolio choice literature (e.g., Heaton and Lucas [2000]).

At the household level, homeownership and other demographic information comes from the 1990 and 1999 Current Population Survey (CPS) March Annual Demographic Supplements, and is matched to the MSA data.¹³ In addition, we impute the probability of a household's staying in the same residence (whether rented or owned) for another year, $P(STAYS)$, as the proportion of households in the same age-occupation-marital status cell (excluding the household in question) that did not move in the previous year.¹⁴

Table I presents summary statistics for the data. Rent risk is quite substantial. Between 1990–1998, the average (across and within MSAs) standard deviation of real rent was 2.9 percent per year, almost half the volatility of real house prices. Sixty percent of the households in our sample own their homes, and they typically pay nearly sixteen times the MSA's annual median apartment rent for them.¹⁵ Much of the variation in these statistics is attributable to cross-sectional differences across the 44

13. We start with more than 110,000 CPS households. Nearly 70,000 are excluded because they do not live in one of our 44 MSAs. Approximately 500 more drop because the household head is under the age of 25, or income or mobility information is missing, yielding a net sample size of 40,274 households.

14. For the imputation, we use the entire CPS sample of households, excluding those with heads under the age of 25 or in the military. We form cells, by year, with households within one of seven 10-year age brackets (25–34, 35–44, 45–54, 55–64, 65–74, 75–84, and 85–94), sixteen occupations (the CPS's "major occupation" code), and seven marital statuses (the CPS marital status definitions). As is customary, we use the age and occupation of the household head.

15. One reason that owner-occupied housing commands a large multiple over rents is that the median house price reflects a greater quantity of (or, equivalently, "nicer") housing than does the average apartment rent. As long as the difference between the amount/quality of housing in the median house and in the average apartment does not spuriously vary across MSAs over time in a way that is correlated with rent variance, it will not affect our estimation.

TABLE I
SUMMARY STATISTICS

Variable	MSA data	Household-level data (CPS)		
	1990–1998	1990 & 1999		
	Mean	Mean	Mean if below median	Mean if above median
Standard deviation (s.d.) of real rent	0.029 [0.017]	0.031 [0.013]	0.021	0.040
S.d. of real house price	0.046 [0.031]			
Real rent growth	0.001 [0.019]			
Real house price growth	0.006 [0.031]			
Average real rent	6,331 [1,505]			
Median real house price	102,773 [49,841]			
Price-to-rent ratio	15.72 [4.08]			
Proportion owning		0.600 [0.490]		
Probability of not moving (imputed $P(STAYS)$)		0.843 [0.116]	0.749	0.936
Probability of not moving × s.d. of real rent		0.026 [0.011]		
Number of observations	396	40,274		

Standard deviations of the variables are in square brackets. Rent growth, house price growth, and the standard deviations of rents and house prices, are all computed over the preceding nine years. The rent data are obtained from Reis. To compute the level of house prices, the MSA median house price from the 1990 Census is inflated to the current year using the Freddie Mac repeat sales price index. Not moving is defined as having not moved in the preceding year. All dollar values are in real (1990) dollars, deflated by the CPI less shelter.

MSAs. For instance, the volatility of rent ranges from 1.4 percent standard deviation in Fort Lauderdale to 7.2 percent in Austin.

For some variables we report separately the means for the top and bottom halves of their distributions, corresponding to how we will group the data in our empirical work. Households who live in “high” rent variance MSAs have a standard deviation of real rent of 4 percent, twice that of those in “low” rent variance MSAs. At the bottom of the table, on average 84 percent of the CPS households did not move in the last year. However, 25 percent (= $1 - 0.749$) of the “mover” households—those below the median in the imputed probability of not moving—moved in the

last year, while only 6.4 percent of the “stayer” households—those above the median—moved.

III. EMPIRICAL METHODOLOGY AND RESULTS

In this section we empirically examine three key implications of our model. First, we estimate the effect of net rent risk on the demand for home owning across MSAs. We examine the effect on both homeownership rates and house prices, since which channel demand operates through depends on the elasticity of supply.¹⁶ While the average effect is theoretically ambiguous in sign, since it depends on the marginal household, using MSA-level data we find that house prices, normalized by rents, do significantly increase with the volatility of local rents.

Second, using household-level data, we find that within MSAs, the probability of home owning increases with a household’s expected horizon in its home. However, the demand for owning can increase with horizon not only because of the rent-risk mechanism, but also because of other unobserved factors correlated with horizon. For instance, demand will also increase with horizon because of fixed moving/transactions costs associated with buying and selling a house that households prefer to amortize over longer stays.

To avoid such factors, we focus on a third implication: whether the probability of owning increases with the interaction of local rent volatility and expected horizon. This interaction effect will isolate the role of net rent risk from other unobserved factors that affect both ownership demand and either horizon or rent volatility separately. We find a significant effect for both the general population, measuring expected horizon by the probability of not moving, and also for the elderly, measuring horizon by age, a proxy for their expected remaining lifetime.

III.A. The Effect of Net Rent Risk on Homeownership Rates

We first test whether net rent risk affects the probability of home owning, both across and within MSAs. Our baseline specification regresses whether a household is a homeowner on local rent volatility (σ_r), a measure of expected horizon (N), and the interaction of the two ($N \times \sigma_r$). Specifically, we estimate probit

16. There is no consensus in the literature regarding the magnitude of the elasticity of supply. See, e.g., Bruce and Holtz-Eakin [1999], Sinai [1998], Ca-pozza, Green, and Hendershott [1996], and Glaeser and Gyourko [2004].

models of the following form using household level data from the 1990 and 1999 CPS:

$$(7) \quad OWN_{i,k,t} = \beta_0 + \beta_1 f(\sigma_r)_{k,t} + \beta_2 g(N)_{i,t} \\ + \beta_3 f(\sigma_r) \times g(N)_{i,k,t} + \theta \mathbf{X}_i + \psi \mathbf{Z}_{k,t} + \zeta_t + \varepsilon_{i,k,t},$$

where i indexes the household, k the MSA it lives in, and t the year. OWN is an indicator variable that takes the value one if the household owns its house and zero otherwise. The standard deviation of rent in market k is denoted by $\sigma_{r,k}$ and is computed over 1980–1989 for the 1990 observations and over 1990–1998 for the 1999 observations. In the reported results, the functions $f(x)$ and $g(x)$ generally take the form of indicator variables that are equal to one when the corresponding variable x , whether rent volatility or horizon, is greater than its sample median. The results are robust to entering the variables linearly. \mathbf{X}_i is a vector of household level controls including log income and dummy variables for race, education, occupation, ten-year age categories, and marital status. $\mathbf{Z}_{k,t}$ is a vector of MSA-level controls including the average rent and median house price in the preceding year, and the average real rent growth and house price growth rates over the preceding nine years. A dummy for 1999 is included to control for the year-specific factor ζ_t .

We use various proxies for the expected horizon N in a house, starting with $P(STAYS)_{i,t}$, the imputed probability that household i does not move during year t . We first investigate the effect of net rent risk on homeownership for the average household in an MSA, by leaving out the interaction term $f(\sigma_r) \times g(N)$ and testing whether the estimated coefficient on rent volatility, β_1 , is positive. The estimated marginal effects are reported in column (1) of Table II. Since a number of the independent variables, including $\sigma_{r,k}$, vary only across markets within a given year, in column (1) we correct the standard errors to account for the correlated shocks within $MSA \times year$ cells.

We find that the average household in a high rent variance MSA is only slightly more likely to be a homeowner than one that lives in a low rent variance MSA, by $\beta_1 = 2.8$ (2.4) percentage points. This statistically insignificant result could mean that the gross rent and asset price risks largely offset each other for the sample households with average expected horizons, in which case there would be little average net effect on either ownership rates or prices. However, below we find that the average effect of net

TABLE II
THE EFFECT OF NET RENT RISK ON THE PROBABILITY OF HOMEOWNERSHIP

	(1)	(2)	(3)	(4)	(5)
Dependent variable: one if household is a homeowner, zero otherwise					
β_1 : σ_r [σ_r = Standard deviation of real rent]	0.028 (0.024)	0.008 (0.022)			
β_2 : N [N = Probability of staying, $P(STAYS)$]	0.036 (0.011)	0.015 (0.013)	0.020 (0.011)		0.018 (0.012)
β_3 : $\sigma_r \times N$ [$N = P(STAYS)$]		0.042 (0.014)	0.029 (0.011)		0.018 (0.013)
β_2 : N [N = Age if over 60]				-0.0006 (0.0007)	
β_3 : $\sigma_r \times N$ [N = Age if over 60]				-0.0029 (0.0014)	
β_4 : r/Y [= Market Rent/ Household Income]					-0.018 (0.016)
β_5 : $r/Y \times N$ [$N = P(STAYS)$]					0.017 (0.019)
β_6 : $\sigma_r \times r/Y$					-0.021 (0.020)
β_7 : $\sigma_r \times r/Y \times N$ [$N = P(STAYS)$]					0.054 (0.025)
MSA controls	Yes	Yes	No	No	No
MSA \times year dummies	No	No	Yes	Yes	Yes
Household controls	Yes	Yes	Yes	Yes	Yes
Number of observations:	40,274	40,274	40,274	9,699	39,468
Pseudo R^2	0.2352	0.2355	0.2498	0.1989	0.2566

This table reports marginal effects from probit regressions of equation (7), with standard errors in parentheses, estimated using household-level data covering 44 MSAs in 1990 and 1999. All specifications include year dummies. MSA controls include median real rent, median real house price, real rent growth, and real house price growth. Household controls include log household income and dummies for the head's occupation, race, education, marital status, and age. In columns (1)–(3) and (5), MSAs are deemed to have high rent variance if σ_r is above the median household's value of 2.8 percent. In column (4), the cutoff is the seventy-fifth percentile household's value of 4.1 percent. The probability of staying is high if the household is above the median probability of 88 percent. All dollar values are in real (1990) dollars, deflated by the CPI less shelter. In columns (1) and (2) the standard errors are adjusted for correlation within MSA/year. Column (5) excludes the outliers with the 1 percent highest and lowest values of MSA average rent to household income.

rent risk is significantly capitalized into house prices, suggesting that the result here instead reflects relatively inelastic supply. Unsurprisingly, households with long expected horizons are more likely to be homeowners. The estimated coefficient β_2 implies that households with the longest horizons, or “stayers,” are 3.6 percentage points more likely to own their homes than are “movers.” This result could, of course, partly reflect transaction

costs or other omitted variables that are correlated with horizon.

The more compelling test of the rent hedging benefit of owning focuses on β_3 and the interaction of rent risk and expected horizon. The difference in the probability of home owning between longer and shorter expected horizon households should increase with the rent variance, so β_3 is expected to be positive. In addition, while unobserved MSA-level characteristics could potentially bias the coefficient β_1 on rent variance alone, β_3 should still be consistent since it depends only on the interaction of household-level characteristics with the MSA-level rent variance. The interaction term is set to one if a household both lives in an MSA with a standard deviation of rent above the median and also is above the median in expected horizon.

In column (2) the estimate of β_3 is indeed positive and statistically significant. Compared with the difference between “stayers” and “movers” in low rent variance MSAs, “stayers” in high rent variance MSAs are 4.2 percentage points more likely to own their home than “movers” in the same places. Relative to the 2003 national average probability of home owning of 68 percent (which increased by only about 3.5 percentage points over the prior twenty years), this is also an economically significant effect. For comparison, based on the estimate of β_1 , the below-median horizon households (with an average horizon of about four years) are less than one percentage point more likely to own their home if they live in a high rent variance MSA.

To control for unobserved MSA characteristics, column (3) includes $MSA \times year$ dummies, at the expense of not identifying purely MSA-level factors such as (uninteracted) $\sigma_{r,k}$. This specification compares homeownership probabilities of long- and short-horizon households *within* each MSA. The estimated coefficient β_3 on the interaction term declines in magnitude to 2.9 (1.1) percentage points, but remains statistically significant.

These results are robust to a number of sensitivity checks, including controlling separately for house price volatility, and measuring $P(STAY)$ as the probability of not moving out of the county, to distinguish moves more likely going to relatively uncorrelated housing markets.¹⁷

17. The reported specifications do not separately control for house price volatility because it is just an endogenous function of rent volatility. Equations (7)

An alternative measure of horizon is, for the elderly, their expected remaining lifetime. As the elderly approach the end of their lives, they expect to face fewer rent shocks, and their asset price risk is closer at hand (and is especially salient if they want to bequeath the monetary proceeds from selling their home at death to their children). According to the model, the decline in the probability of homeownership with age should be steepest in high rent variance MSAs. Other mechanisms affecting the age-ownership profile are unlikely to be systematically different in high and low rent volatility MSAs.

This effect of rent volatility can be seen directly in the unconditional homeownership rates by age. Using the pooled 1990 and 1999 CPS cross sections, we divided our 44 MSAs into high- and low-variance markets depending upon whether they were in the top quartile of rent variance or below the median, and used a kernel regression to compute the homeownership rate by age in both sets of markets. The result appears in Figure I. The probability of homeownership is generally larger in high rent-variance MSAs than in low variance MSAs, and declines with age starting around age 65. The probability does indeed fall faster for households in high rent variance MSAs.¹⁸

We repeat the analysis of column (3) now for the elderly, limiting the sample to households over the age of 60, and using the age of the household's head as an inverse proxy for expected horizon, with a high age corresponding to a shorter horizon. To more sharply identify the effect of rent risk given the smaller sample, $f(\sigma_r)$ is equal to one if the MSA is in the top quartile of rent variance and zero otherwise.

The results, reported in column (4), are consistent with the rent hedging mechanism. In particular, the coefficient β_3 on the interaction term is significantly negative, implying that homeownership declines more rapidly with age in high rent variance

and (8) appropriately capture the net effect of both risks. The results also are robust to measuring N as $1/(1 - P(STAYS))$, to interacting $P(STAYS)$ with all the observable MSA characteristics (and the year dummy), and to controlling for the correlation of rents and income. We do not focus on the across-county probability of moving because the CPS does not identify which county a household moves to; rather, only whether the household came from outside the current county. Additional empirical specifications are reported in the working paper version of this paper.

18. The difference in homeownership probabilities between high and low variance MSAs peaks at around age 60, perhaps because the elderly are more risk-averse than younger households (a level effect). We focus on how these probabilities decline with remaining lifetime (a slope effect) for the elderly; for younger households, remaining lifetime and horizon do not coincide as closely.




FIGURE I
Kernel-Smoothed Age Profile of Homeownership, by Rent Variance

MSAs. Relative to people over 60 in low rent variance MSAs, the probability of homeownership for people over 60 in high rent variance MSAs falls by 0.29 (0.14) percentage points more per year of age.

In addition to increasing with horizon, a household's effective net rent risk is likely to increase with the importance of housing costs in its budget. We measure this by the ratio of the average rent in the MSA to actual household income, r/Y . We use MSA-level rent instead of the household's own rent since the former is exogenous to the household and is defined even for homeowners. Households for whom housing is a larger portion of their budget might be more sensitive to net rent risk since they are implicitly taking a larger gamble. If so, homeownership rates should be highest among those with high rent-to-income ratios, long expected horizons, and high rent volatility (i.e., interacting all three terms, $r/Y \times N \times \sigma_r$). We test this "budget share" hypothesis by dividing the sample based on r/Y , and fully interacting it with the variables of interest. High rent variance and "stayer" households are again defined relative to their respective medians, and high rent-to-income households are those in the top quartile.

In column (5) we find that the coefficient on the triple-interaction term is positive and significant: being a long-horizon

household, with a large budget share, in a high rent variance MSA, raises the probability of homeownership by 5.4 percentage points relative to other households. This specification controls for unobservable differences between movers and stayers, high rent variance MSAs and low variance MSAs, and high rent-to-income and low rent-to-income households, as well as the corresponding pairs of binary interactions.

III.B. The Effect of Net Rent Risk on the Price-to-Rent Ratio

Turning to house prices, we examine the effect of net rent risk on the price-to-rent ratio. Using this ratio controls for shocks to the overall housing market, which impact both owner-occupied and rental housing. In markets where rent volatility is greater (assuming the supply of owner-occupied housing is not fully elastic), we would expect to see a larger price-to-rent ratio, as prices capitalize the additional value of the rent hedging benefit of owning above and beyond the expected present value of the housing service flow.

We estimate the following equation using OLS on the MSA-level panel data for 1990–1998:

$$(8) \quad (P/r)_{k,t} = \alpha_0 + \alpha_1(\sigma_r)_{k,t} + \psi Z_{k,t} + \zeta_t + u_{k,t},$$

where $(P/r)_{k,t}$ is the price-to-rent ratio in MSA k in year t , $\sigma_{r,k,t}$ is its standard deviation of rent, and $Z_{k,t}$ is its growth rate of real rent. According to the model, P/r should be higher for MSAs with higher expected future rent growth, so ψ is expected to be positive. The coefficient α_1 will then capture the net rent risk premium. The year dummies ζ_t control for differences over time that are common to all MSAs.

In Table III, house price-to-rent ratios appear to incorporate both expected future rents and the associated risk premiums. In particular, MSAs with more volatile rents have significantly greater price-to-rent ratios. As reported in column (1), the estimated coefficient $\alpha_1 = 34.5$ (11.9) is positive and significant. The estimated ψ is also positive and significant, at 69.0 (14.7). These price results are consistent with our model and more generally with other asset-pricing models. The last two rows of Table III help gauge the economic significance of these results. A one standard deviation increase in $\sigma_{r,k,t}$ is estimated to increase the price-to-rent ratio by 0.62. Since the mean price-to-rent ratio is 15.7, this amounts to a 3.9 percent rise in house prices, holding rents constant.

TABLE III
THE EFFECT OF NET RENT RISK ON THE PRICE-TO-RENT RATIO

	(1)	(2)	(3)
Dependent variable: Price-to-rent ratio			
α_1 : Standard deviation of real rent (σ_r)	34.52 (11.88)	11.04 (5.55)	10.10 (3.81)
ψ : Real rent growth	68.99 (14.68)	16.73 (4.67)	18.14 (5.23)
Controls for MSA fixed effects?	No	MSA dummies	First differences
Number of observations	396	396	352
R^2	0.0486	0.9471	0.1609
A one s.d. increase in σ_r leads to a . . . increase in the price-to-rent ratio	0.62 (0.21)	0.20 (0.10)	0.18 (0.07)
A one s.d. increase in σ_r leads to a . . . percent increase in house prices, holding rent constant	3.9	1.3	1.1

Estimation is by OLS, following equation (8). Standard errors are in parentheses. Number of observations equals 44 MSAs per year over the 1990–1998 time period. All specifications include year dummies. σ_r and rent growth rates are computed based on the previous (rolling) nine years. A one standard deviation increase in σ_r is 0.018 (from a mean of 0.031). The average price-to-rent ratio is 15.72.

Column (2) incorporates MSA dummies, and so uses the within-MSA variation in rent volatility, rent growth, and the price-to-rent ratio over time to identify the rent hedging mechanism. (Recall that the rent variance $\sigma_{r,k,t}$ and (rent growth) $_{k,t}$ within an MSA change over time as the rolling window over which we compute them moves.) The results are qualitatively similar to the previous ones though smaller in magnitude, which is not surprising considering that only within-MSA variation is being used for identification. A one standard deviation increase in σ_r now increases house prices by 1.3 percent, given rents. Column (3) accounts for MSA level heterogeneity by instead estimating equation (8) in first differences. This specification emphasizes new information that arrives over time, since the difference in the computed rent variance between one year and the previous year is due to adding the most recent year of data and discarding the oldest year in the rolling window used to calculate σ_r . The results are very similar to those in column (2). In all columns these conclusions are robust to controlling additionally for the volatility of house prices.

Overall, these results show that at least some of the rent

hedging benefits of homeownership are capitalized into local price-to-rent ratios, even for the average household.

IV. CONCLUSION

Since all households start life with an implicit short position in housing, they are all exposed to housing market risk. Renters face substantial fluctuations in their rents, whereas owners are hedged against uncertainty in spot housing costs. But owners in turn face asset price risk when they move (or die). In a simple tenure choice model with endogenous house prices, we show how the demand for owning trades off these two risks. The sign of the net risk depends on a household's expected horizon in its home, and the magnitude increases multiplicatively with the volatility of rents. For households with long horizons, the rent risk is more likely to dominate the asset price risk, increasing the demand for owning, as the number of rent risks increases and the asset price risk arrives further in the future. Greater spatial correlation in housing costs when a household moves defers the asset price risk further, effectively extending the household's horizon. For these reasons, the asset price risk of owning can be small even when considered in isolation. In addition, because the value of a housing asset is inherently tied to the size of the corresponding implicit rent liability, their offsetting effects reduce the overall wealth effects from changes in housing costs.

Empirically, we confirmed that the probability of homeownership, as well as the multiple of rents households are willing to pay to own their homes, increases with net rent risk. Further, the probability of owning significantly increases with the interaction of rent risk and horizon. That is, the demand for owning increases faster with rent risk for households with longer horizons. Overall, the rent hedging benefit of owning appears to be a significant factor in explaining homeownership.

Contrary to conventional wisdom, rent risk can dominate the asset price risk, especially for households with long horizons or when housing costs are spatially correlated. In that case, greater housing market volatility can actually increase the demand for owning. To gauge the significance of these results, consider the effect of eliminating rent volatility altogether, using estimates of a variant of equation (7) in which the covariates enter linearly. For the 75 percent of sample households with the longest ex-

pected horizons, the probability of homeownership would decline by 3.3 percent, as the net rent risk avoided by owning is eliminated. For the remaining households, who have short horizons, the probability of homeownership would increase by as much as 10 percent if rent volatility were eliminated, since for them the asset price risk dominates the rent risk. The effect of rent risk on house prices is also large. Using the smaller estimates from column (2) of Table III, if there were no rent risk, then house prices relative to rents would decline by 2.3 percent on average and by as much as 7 percent in some MSAs.

APPENDIX

We present the general case of the components of RP from equation (1), allowing rents to be persistent ($\varphi > 0$) and spatially correlated ($\rho > 0$). The first component is the rent risk premium:

$$\begin{aligned}
 (A1) \quad \pi_R \approx \frac{\alpha}{2} \kappa^2 & \left\{ \sigma_A^2 \left(\sum_{n=1}^{N-1} \left[\delta^n + \sum_{i=n+1}^{N-1} \delta^i \varphi^{i-n} + \rho \sum_{i=N}^{2N-1} \delta^i \varphi^{i-n} \right]^2 \right. \right. \\
 & \left. \left. + \sum_{n=N}^{2N-1} \rho^2 \left[\delta^n + \sum_{i=n+1}^{2N-1} \delta^i \varphi^{i-n} \right]^2 \right) \right. \\
 & \left. + \sigma_B^2 \left(\sum_{n=1}^{N-1} \left[\rho \left(\delta^n + \sum_{i=n+1}^{N-1} \delta^i \varphi^{i-n} \right) + \sum_{i=N}^{2N-1} \delta^i \varphi^{i-n} \right]^2 \right. \right. \\
 & \left. \left. + \sum_{n=N}^{2N-1} \left[\delta^n + \sum_{i=n+1}^{2N-1} \delta^i \varphi^{i-n} \right]^2 \right) \right\}.
 \end{aligned}$$

Relative to the special case in equation (3), in the first square brackets multiplying σ_A^2 , the added second term reflects the fact that if $\varphi > 0$ then each early location-A shock (η_n^A) continues to affect the subsequent rents incurred while living in A through year $N - 1$. The added third term reflects the fact that if also $\rho > 0$, then each early location-A shock continues to affect the late rents r_N^B to r_{2N-1}^B incurred while living in B. The second outer summation term following σ_A^2 (outside the second square brackets) measures the effect of the late location-A shocks (η_N^A to η_{2N-1}^A) on the late rents incurred in B, if $\rho > 0$. The terms multiplying σ_B^2 are analogous.

The risk premium for owning takes the form,

$$\begin{aligned}
 \text{(A2)} \quad \pi_O &\approx \frac{\alpha}{2} \kappa^2 \left(\frac{\delta^N}{1 - \delta\varphi} \right)^2 \\
 &\times \left\{ \sigma_A^2 \left([1 - \rho(1 - \delta^N \varphi^N)]^2 \sum_{i=1}^N \varphi^{2(N-i)} + (\rho\delta^N)^2 \sum_{i=N+1}^{2N} \varphi^{2(2N-i)} \right) \right. \\
 &\quad \left. + \sigma_B^2 \left([\rho - (1 - \delta^N \varphi^N)]^2 \sum_{i=1}^N \varphi^{2(N-i)} + (\delta^N)^2 \sum_{i=N+1}^{2N} \varphi^{2(2N-i)} \right) \right\}.
 \end{aligned}$$

Relative to the special case in equation (6), the first summation term multiplying σ_A^2 captures the effects of the early shocks η_1^A to η_N^A on all three prices P_N^A , P_N^B and P_{2N}^B . Similarly, the second summation captures the effects of the late shocks η_{N+1}^A to η_{2N}^A on P_{2N}^B . In the IID case with $\varphi = 0$, the price risks are of the same magnitude as the corresponding individual rent risks. But as φ increases, more of the prior rent shocks accumulate and get embedded into the corresponding prices, increasing the magnitude of the summation terms in equation (A2). The $(1 - \delta^N \varphi^N)$ terms in (A2) reflect the intertemporal natural hedge. While the early shocks η_1^B to η_N^B (and if $\rho > 0$, η_1^A to η_N^A) lead to uncertainty in the future purchase price P_N^B , if $\varphi > 0$ this uncertainty is partially offset by their oppositely signed effect on the subsequent sale price P_{2N}^B of the same house.

Similarly, in the net risk premium in P_N^B that is indirectly capitalized into P_0^A , the additional terms relative to equation (5) reflect the effect of persistence in augmenting the total amount of rent shocks and the magnitude of the sale risk:

$$\begin{aligned}
 \text{(A3)} \quad \delta^N (\pi_R^B - \pi_O^B) &\approx \delta^N \frac{\alpha}{2} \kappa^2 (\sigma_B^2 + \rho^2 \sigma_A^2) \\
 &\times \left\{ \sum_{n=1}^{N-1} \left(\delta^n + \sum_{i=n+1}^{N-1} \delta^i \varphi^{i-n} \right)^2 - \left(\frac{\delta^N}{1 - \delta\varphi} \right)^2 \sum_{i=0}^{N-1} \varphi^{2i} \right\}.
 \end{aligned}$$

THE WHARTON SCHOOL, UNIVERSITY OF PENNSYLVANIA, AND NATIONAL BUREAU OF ECONOMIC RESEARCH

REFERENCES

Bajari, Patrick, C. Lanier Benkard, and John Krainer, "House Prices and Consumer Welfare," mimeo, May 2003.

- Ben-Shahar, Danny, "On the Optimality of the Hybrid Tenure Mode," *Journal of Housing Economics*, VII (1998), 69–92.
- Berkovec, James, and Don Fullerton, "A General Equilibrium Model of Housing, Taxes, and Portfolio Choice," *Journal of Political Economy*, C (1992), 390–429.
- Bruce, Donald, and Douglas Holtz-Eakin, "Fundamental Tax Reform and Residential Housing Values," *Journal of Housing Economics*, VIII (1999), 249–271.
- Brueckner, Jan, "Consumption and Investment Motives and the Portfolio Choices of Homeowners," *Journal of Real Estate Finance and Economics*, XV (1997), 159–180.
- Campbell, John, and Joao Cocco, "Household Risk Management and Optimal Mortgage Choice," *Quarterly Journal of Economics*, CXVIII (2003), 1449–1494.
- Capozza, Dennis R., Richard Green, and Patrick H. Hendershott, "Taxes, Home Mortgage Borrowing and Residential Land Prices," in Aaron and Gale, eds., *Fundamental Tax Reform* (Washington, DC: The Brookings Institution, 1996), pp. 181–204.
- Case, Karl, John Quigley, and Robert Shiller, "Comparing Wealth Effects: The Stock Market versus the Housing Market," mimeo, University of California, Berkeley, May 2003.
- Case, Karl, and Robert Shiller, "The Efficiency of the Market for Single-Family Homes," *American Economic Review*, LXXIX (1989), 125–137.
- Chan, Sewin, "Spatial Lock-In: Do Falling House Prices Constrain Residential Mobility?" *Journal of Urban Economics*, XLIX (2001), 567–586.
- Cocco, Joao, "Hedging House Price Risk with Incomplete Markets," mimeo, London Business School, September 2000.
- Davidoff, Thomas, "Labor Income, Housing Prices, and Homeownership," mimeo, University of California, Berkeley, April 2003.
- Engelhardt, Gary, "Consumption, Down Payments, and Liquidity Constraints," *Journal of Money, Credit and Banking*, XXVIII (1996), 255–271.
- Flavin, Marjorie, and Takashi Yamashita, "Owner-Occupied Housing and the Composition of the Household Portfolio over the Life Cycle," *American Economic Review*, XCII (2002), 345–362.
- Flavin, Marjorie, and Shinobu Nakagawa, "A Model of Housing in the Presence of Adjustment Costs: A Structural Interpretation of Habit Persistence," mimeo, October 2003.
- Fratantoni, Michael, "Housing Wealth, Precautionary Saving, and the Equity Premium," mimeo, Fannie Mae, July 1997.
- Genesove, David, "The Nominal Rigidity of Apartment Rents," NBER Working Paper No. 7137, May 1999.
- Genesove, David, and Christopher Mayer, "Equity and Time to Sale in the Real Estate Market," *American Economic Review*, LXXXVII (1997), 255–269.
- Glaeser, Edward, and Joseph Gyourko, "Urban Decline and Durable Housing," Wharton Real Estate Center Working Paper No. 382, August 2004.
- Goetzmann, William, "The Single Family Home in the Investment Portfolio," *Journal of Real Estate Finance and Economics*, VI (1993), 201–222.
- Gross, David, and Nicholas S. Souleles, "Do Liquidity Constraints and Interest Rates Matter for Consumer Behavior? Evidence from Credit Card Data," *Quarterly Journal of Economics*, CXVII (2002), 149–185.
- Gyourko, Joseph, and Joseph Tracy, "Using Home Maintenance and Repairs to Smooth Variable Earnings," Wharton Real Estate Center Working Paper No. 432, March 2004.
- Haurin, Donald, "Income Variability, Homeownership, and Housing Demand," *Journal of Housing Economics*, I (1991), 60–74.
- Heaton, John, and Deborah Lucas, "Portfolio Choice and Asset Prices; The Importance of Entrepreneurial Risk," *Journal of Finance*, LV (2000), 1163–1198.
- Hendershott, Patrick, and Joel Slemrod, "Taxes and the User Cost of Capital for Owner-Occupied Housing," *AREUEA Journal*, X (1983), 375–393.
- Henderson, J. Vernon, and Yannis M. Ioannides, "A Model of Housing Tenure Choice," *American Economic Review*, LXXIII (1983), 98–113.
- Jappelli, Tullio, Jorn-Steffen Pischke, and Nicholas Souleles, "Testing for Liquidity Constraints in Euler Equations with Complementary Data Sources," *Review of Economics and Statistics*, LXXX (1998), 251–262.

- Ortalo-Magné, François, and Sven Rady, "Tenure Choice and the Riskiness of Non-housing Consumption," *Journal of Housing Economics*, XI (2002a), 266–279.
- Ortalo-Magné, François, and Sven Rady, "Homeownership: Low Household Mobility, Volatile Housing Prices, High Income Dispersion," University of Wisconsin Working Paper, November 2002b.
- Poterba, James, "Tax Subsidies to Owner-Occupied Housing: An Asset Market Approach," *Quarterly Journal of Economics*, XCIX (1984), 729–752.
- Poterba, James, and Andrew Samwick, "Household Portfolio Allocation over the Life Cycle," NBER Working Paper No. 6185, September 1997.
- Rosen, Harvey, "Housing Decisions and the U. S. Income Tax," *Journal of Public Economics*, XI (1979), 1–23.
- Rosen, Harvey, Kenneth Rosen, and Douglas Holtz-Eakin, "Housing Tenure, Uncertainty, and Taxation," *Review of Economics and Statistics*, LXVI (1984), 405–416.
- Sinai, Todd, "Are Tax Reforms Capitalized into House Prices?" mimeo, The Wharton School, December 1998.
- Skinner, Jonathan, "Housing Wealth and Aggregate Saving," *Regional Science and Urban Economics*, XIX (1989), 305–324.
- , "Is Housing Wealth a Sideshow?" in David Wise, ed., *Advances in the Economics of Aging* (Chicago: University of Chicago Press for NBER, 1996), pp. 241–268.
- Stein, Jeremy, "Prices and Trading Volume in the Housing Market: A Model with Down-Payment Constraints," *Quarterly Journal of Economics*, CX (1995), 379–406.
- Summers, Lawrence, "Inflation, the Stock Market, and Owner-Occupied Housing," *American Economic Review*, LXXI (1981), 429–434.
- Zeldes, Stephen P., "Consumption and Liquidity Constraints: An Empirical Investigation," *Journal of Political Economy*, XCVII (1989), 305–346.