

The Effects of Price Risk on Housing Demand: Empirical Evidence from U.S. Markets¹

Lu Han

Rotman School of Management, University of Toronto, Ontario, Canada M5S 3E6

¹This paper is a substantially revised version of the first chapter of my Stanford University dissertation. I thank Takeshi Amemiya, Susan Athey, Pat Bajari, Tim Bresnahan, Liran Einav, Matt Kahn, Rob McMillan, John Shoven, and Will Strange for very helpful discussions and suggestions. I am also grateful to the editor, Matthew Spiegel, and to an anonymous referee for comments that greatly improved the paper. This work was funded by the Kapnick Dissertation Fellowship at Stanford University and the Connaught New Staff Matching Grant at the University of Toronto. Email: lu.han@rotman.utoronto.ca. Phone: 416-946-5294. Fax: 416-978-7030.

Abstract

This paper examines how price risk affects housing demand. It identifies two relevant channels: a financial risk effect that reduces demand, and a hedging effect that increases demand since current homes may hedge future housing costs. The latter dominates when hedging incentives are strong, namely when the likelihood of moving up the housing ladder is high and the tendency to move across markets is low. For households with weak hedging incentives, the paper finds negative effects of price risk on the timing and size of home purchases, but positive effects for households with strong hedging incentives.

Keywords: Housing demand; Price risk; Hedging; Self-insurance; Mobility

Purchasing a home is the single most important financial decision that the typical household makes. Housing accounts for a significant fraction of household wealth – around two-thirds for the median household in the United States (Tracy and Schneider (2001)). There is also considerable risk associated with home purchases due to house price volatility. This has been all too apparent in the current housing crisis, with real house prices rising by over 70% between 2001 and 2006, yet falling by over 30% between 2006 and 2009.¹ In light of such dramatic changes, house price risk has recently ranked as one of the most significant risks that U.S. homeowners face (Berson and Henry (2008)), attracting keen interest from economists and policy makers alike.

This paper focuses on estimating the effects of price risk on housing demand. Conceptually speaking, an increase in price risk affects housing demand through two channels. The first is the financial risk effect associated with a household’s current dwelling. Under standard assumptions, an increase in asset price uncertainty increases the financial risk associated with future asset returns, and hence makes the asset less desirable. This negative effect on asset demand tends to be particularly large for housing because house price risk is not readily diversifiable (Caplin, Chan, Freeman, and Tracy (1997)).² Furthermore, the large and lumpy transaction costs involved in each home sale make it almost impossible for households to make small adjustments in response to price fluctuations (Haurin and Gill (2002)). Thus, conventional wisdom suggests that increasing house price risk will reduce the demand for housing.

The second channel is the hedging effect associated with future housing consumption. For the typical U.S. household, housing is not only a financial asset but also a necessary consumption good (Henderson and Ioannides (1983), Brueckner (1997)); and a key feature of home ownership is that it gives a household the ability to use an early purchase to hedge against future housing cost risk (Berkovec and Fullerton (1992), Ortalo-Magne and Rady (2002)). For example, if the price of a future house is positively correlated with the price of the current house, the household may make an earlier or larger home purchase in order to offset future housing cost risk — often referred to, in real estate parlance, as “getting into the market.”

The goal of this paper is to test for the presence of these two alternative but not necessarily

mutually exclusive forces — the financial risk effect and the hedging effect — on housing demand. Existing empirical work focuses primarily on the choice between renting and owning and finds mixed evidence in that context. For example, findings in Rosen, Rosen, and Holtz-Eakin (1984) point to the operation of the financial risk effect, while Banks, Blundell, Oldfield, and Smith (2004) and Sinai and Souleles (2005) provide support for the importance of hedging incentives.

Building on previous work, this paper is the first to estimate the effects of price risk on the timing and size of home purchases made by *existing* homeowners.³ This focus is motivated by two observations. First, existing homeowners account for about 70% of U.S. households; and compared with first-time home buyers, the hedging demand by existing homeowners is less likely to be affected by liquidity constraints. Thus, a focus on existing homeowners should yield particularly powerful tests for, and useful independent evidence regarding, the effects of price risk on housing demand. Second, as noted by Goetzmann and Spiegel (2002), the traditional measurement of tenure choice in the context of a static home ownership likelihood obscures some important issues inherent in the dynamic home purchase decision. In particular, for existing homeowners, their current housing demand depends on their future housing consumption plans, as well as transaction costs associated with selling previous homes.

Capturing these aspects for existing homeowners requires a demand model that not only highlights the sequential nature of home purchase decisions made at different stages of the life-cycle, but also the simultaneity between the timing and size of home purchase decisions. From a methodological perspective, this paper extends the standard type-2 Tobit model in a cross-sectional context (Amemiya (1985)), yielding a dynamic correlated random-effect housing demand model. This formulation is rich enough to meet both requirements just described, yet simple enough to permit a separation of the financial risk and hedging motives in testing the effects of price risk on housing demand.

As a precursor to estimating the model, the paper proxies home purchase size by imputing a hedonic-adjusted measure of housing quantity that is comparable both over time and across markets. To proxy hedging incentives, the paper constructs a household-specific time-varying measure that captures the probability that a household will trade up to a bigger house in the next five years

within the same market. Given that it is not possible to fully characterize a homeowner's housing consumption plans, there will invariably be a component in the hedging incentive that is unobserved to the econometrician; and to the extent that such unobservables also affect housing demand, an endogeneity problem arises and may bias the demand estimates. To correct this potential endogeneity problem, the paper develops an instrumental variables strategy based on whether the household head's occupation is licensed, the rationale being that having a licensed occupation constrains the household's geographical mobility and hence strengthens its hedging incentives while not affecting housing demand directly.

For estimation purposes, the paper draws on two national panel datasets. The primary data source is the Panel Study of Income Dynamics (PSID), which contains household-level information on home purchases and demographics over the period 1980-1997. These micro-level panel data are useful for examining intertemporal dynamics in home purchase behavior. More importantly, panel data allow one to control for time-invariant unobserved differences in household characteristics that may confound estimation. Quarterly market-level housing price information for 1980-1997 comes from a secondary dataset, the Conventional Mortgage Home Price Index (CMHPI), which provides rich metropolitan-level variation in house prices over time.

Estimates of the model provide evidence that both the financial risk effect and the hedging effect operate strongly. For households with weaker hedging incentives (a hedging propensity in the bottom third of the distribution), a one percentage-point increase in house price risk, measured by the standard deviation of the return, *decreases* the probability of making a transaction by 0.68%. Conditional on transacting, the size of the home purchase also *decreases* by 1.42%. For households with stronger hedging incentives (a hedging propensity in the top third of the distribution), a one percentage-point increase in house price risk *increases* the probability of making a transaction by 0.45%; and conditional on transacting, the size of a home purchase by 1.06%.

These effects are quite robust. They persist when using alternative measures of house size, such as deflated house value, imputed square footage and the number of rooms. Auxiliary tests show that other plausible sources of bias, such as age-varying risk aversion, are not responsible for the main results.

Given the general consistency across specifications, the paper’s findings indicate that households are not only aware of house price risk but also take advantage of the intertemporal relationship between sequential home purchases to reduce this risk. Thus, even in the presence of incomplete markets, where families lack formal hedging instruments, households are able to use self-hedging mechanisms to deal with housing cost risk. This means that the potential gains from recent market developments in housing futures may be limited, as they could crowd out households’ self-hedging efforts.

The findings not only explain the ability of households to self-insure against house price risk, but also draw attention to the difference in hedging behavior across households and across different stages of the life-cycle for the same household. For example, younger households have stronger hedging incentives when compared with older households. And for those in the same age bracket, having a household head with a licensed occupation and school-age children also increases households’ hedging propensity by constraining their geographical mobility.

The remainder of the paper is organized as follows. Section 1 presents theoretical foundations for the empirical tests in the paper. Section 2 lays out the empirical specification for modeling intertemporal housing demand for existing homeowners, and Section 3 describes the data and constructs the key variables. Section 4 presents identification strategies. Section 5 then discusses results from several alternative specifications. Section 6 provides necessary robustness checks to confirm the validity of the main results, and Section 7 concludes.

1 Price Risk and Housing

This section will discuss the theoretical foundations of the estimation approach in a heuristic way, since the paper is primarily empirical in its focus. A detailed stochastic life-cycle model of home purchase behavior is presented in a companion paper (Han (2008)).

The key question addressed in this paper is: How does an increase in house price risk affect existing homeowners’ housing demand? Families that already own a home often have reasons to move and then purchase a new home (Stein (1995)). This recurring need to transact is the driving force

behind hedging incentives for existing homeowners. Consider a household that has no incentive to hedge against future home purchase risk. In a volatile housing market, such a household would take a smaller housing position to avoid the financial risk associated with future housing wealth. Thus for households of this type, housing demand is a *decreasing* function of price risk. Now consider a second household that is identical to the previous one except that this household has a much stronger hedging incentive. In a volatile market, a larger housing position today, although risky in terms of returns, actually reduces the risk of any future home purchase. For households with sufficiently strong hedging incentives, the financial risk effect is dominated by the hedging effect. In this case, housing demand is an *increasing* function of price risk.

The implication that follows from this line of reasoning is that price risk has two effects that work in opposite directions: a financial risk effect and a hedging effect. The net effect of price risk on demand depends crucially on a household's hedging incentives. To test the strength of these forces empirically, we would like a time-varying measure of hedging incentives at the household level. Simulations in Han (2008) show that the strength of homeowners' hedging incentives is governed by two key factors in their housing consumption plans: (1) the spatial correlation in prices across current and future desired houses and, (2) the size of the future desired house relative to the current house. Intuitively, for homeowners moving across uncorrelated markets, the current housing investment provides no hedging value against future home purchase risk; for homeowners moving down the housing ladder, existing homes provide more than a sufficient hedge. In both cases, there are no incentives to purchase an additional amount of housing for hedging purposes. Thus, theory generates a clear characterization of households' hedging incentives that motivates the empirical measure of a hedging propensity score.

Theory also guides the empirical specification in other ways. First, life-cycle models have the feature that any housing decision made today dictates housing opportunities in the future, and thus is dependent on the future housing consumption plan.⁴ Empirically, this requires a dynamic demand model in which housing decisions are not estimated as a one-time decision but rather as a series of decisions made in each period over the life-cycle. Second, as a key characterization of existing homeowners, lumpy transaction costs imply that timing and size of home purchases should be estimated simultaneously. The optimal decision rules derived from consumption durable literature (e.g., Gross-

man and Laroque (1990)) generate useful exclusion restrictions for identifying the joint demand model: for any household in any given period, housing stock imbalance and lumpy transaction costs affect the timing but not the size of home purchases.⁵

Before proceeding, I should highlight one caveat about the mapping from the theory to the empirical work. In practice, for households that are financially constrained from making a new home purchase, remodeling provides an alternative way to hedge against future housing cost risk. Unfortunately, the data do not allow me to observe remodeling expenditures. Since remodeling allows one to increase the current housing position without incurring additional transaction costs, the overall degree of self-hedging among existing homeowners could be even greater than the hedging effect estimated in this paper.

2 Empirical Model

2.1 Model specification

This section presents an empirical specification of existing homeowners' housing demand over the life-cycle. Figure 1 shows that household i makes J_i transactions in T_i periods ($J_i \leq T_i$). In each period t , household i decides whether to transact or to stay and, if transacting, chooses the size of the home purchase. The econometric framework is laid out in three steps. The first is a specification of household i 's decisions at time t on whether to transact and how much to purchase if transacting. The second step is to model the observed sequence of home purchase decisions made by household i in *all periods* for $t = \{1, \dots, T_i\}$. The final step involves constructing the likelihood of observing sequences of housing decisions for *all households* in the sample.

I first begin with the specification of household i 's decisions at time t in its j th purchase spell, where a purchase spell is defined as the time interval between two sequential home purchases. Let $d_{it} \in \{0, 1\}$ indicate household i 's decision about whether to make a transaction in period t . If d_{it} is zero, the household stays in the current house; if not, the size of the j th home purchase is denoted by H_{it}^j . The joint home purchase decisions, d_{it} and H_{it}^j , are assumed to be functions of price risk,

VOL_{it} , hedging propensity, HP_{it} , household i 's demographics and market conditions, x_{it} , instruments for the timing decision, z_{it} , and disturbance terms, ϵ_{it} and e_{it} . Thus, the housing demand specification is given by:

$$d_{it} = 1 \text{ if } \beta_0 + \beta_1 VOL_{it} + \beta_2 HP_{it} \times VOL_{it} + \beta_3 HP_{it} + x_{it}\beta_4 + z_{it}\beta_5 + \epsilon_{it} \geq 0. \quad (1)$$

$$H_{it}^j = \gamma_0 + \gamma_1 VOL_{it} + \gamma_2 HP_{it} \times VOL_{it} + \gamma_3 HP_{it} + x_{it}\gamma_4 + e_{it} \text{ if } d_{it} = 1. \quad (2)$$

For household i , the j th home purchase occurs at time t if the time- t transaction benefits exceed the transaction costs. This is captured by the linear inequality expression (1). Conditional on transacting at time t , the size of the j th home purchase is determined by the right hand-side of equation (2).

The key variables of interest are VOL_{it} and $HP_{it} \times VOL_{it}$. VOL_{it} is the conditional forecast of house price risk at time t in the market where household i resides. γ_1 and β_1 therefore capture the financial risk effect. The hedging effect is explored by including an interaction term between price volatility VOL_{it} and hedging propensity score HP_{it} , both of which are constructed in Section 3. Testing the theory is formally equivalent to testing whether $\beta_1 < 0$, $\beta_2 > 0$, $\gamma_1 < 0$, and $\gamma_2 > 0$.

HP_{it} is included to control for the impact of the hedging propensity alone. The control vector x_{it} contains household i 's demographics, the housing user cost, and market conditions, all evaluated at time t . The vector z_{it} contains the necessary instrumental variables that affect the timing equation but not the size equation. Variables included in x_{it} and z_{it} are discussed in Section 3.

Finally, to account for the possibility that the disturbances in the housing demand specification contain both persistent and transitory components, the model assumes that the distribution of (ϵ_{it}, e_{it}) is characterized by the error component structure:

$$\epsilon_{it} = \mu_i + \eta_{it}, \quad (3)$$

$$e_{it} = \alpha_i + \nu_{it}, \quad (4)$$

where μ_i and α_i are household-specific permanent components with covariance matrix $\Sigma_{\mu\alpha}$, and $\Sigma_{\mu\alpha}$ consists of elements $\sigma_{\mu\mu}$, $\sigma_{\alpha\mu}$, and $\sigma_{\alpha\alpha}$. They can be viewed as reflecting permanent unobserved

differences in housing tastes across households. η_{it} and ν_{it} are independently and identically distributed transitory components (across both i and t) that have zero means and a covariance matrix $\Sigma_{\eta\nu}$, where $\Sigma_{\eta\nu}$ consists of elements $\sigma_{\eta\eta}$, $\sigma_{\eta\nu}$, and $\sigma_{\nu\nu}$. They are intended to capture time-varying unobserved shocks that affect home purchase decisions.

The above empirical specification is a dynamic type-2 Tobit model. The original type-2 Tobit model was developed in Amemiya (1985) in a one-period setting. Unlike Amemiya (1985), the present model makes use of panel data and estimates the sequence of home purchase decisions made over the life-cycle. Such modeling is particularly useful for examining intertemporal home purchase behavior for existing homeowners for three reasons. First, the permanent effects pick up time-invariant unobservables common at the household level; this allows one to separate the intertemporal dependence in home purchase decisions due to transaction cost effects or hedging effects from “spurious” state dependence, which is instead due to the serial correlation generated by the presence of household permanent effects. Second, the model can be easily extended to incorporate a correlated random effects specification, in which household permanent effects are correlated with the time-varying factors. This allows one to address some endogeneity concerns associated with using expected future variables in a housing demand model. The details of extending the current model to a correlated random effects model are presented in Appendix 1.

Third, and more importantly, the type-2 Tobit model allows one to address the sample selection problem in the housing demand model. For example, if the net effect of price risk on housing demand is negative, then households facing little risk will, on average, have larger benefits from transacting and therefore a higher transaction probability than those facing greater price risk. If the timing equation was not estimated, one would only observe the purchases made by those facing low price risk who had comparatively larger housing demand. In this case, the estimates of price risk effects would suffer from selection bias.⁶

To address this selection problem, I combine econometrical modeling techniques with exclusion restrictions motivated by transaction cost theories. By integrating the timing and size decision into a dynamic type-2 Tobit framework, the model allows selection to enter through both permanent

and transitory levels. Under normal distribution assumptions, selection terms can be derived that have known functional forms: the significance of these terms can then be tested in a standard way. In addition, the model also exploits an exclusion restriction strategy where z_{it} includes two sets of variables: transaction costs and housing stock imbalances between the currently desired stock and the existing stock inherited from the previous purchase decision. While transaction costs impose barriers, housing stock imbalances represent disequilibrium motivations for the next transaction. Once a transaction occurs, variables pertaining to the past are sunk factors and do not affect current home purchase size. This makes transaction costs and housing stock imbalance suitable instruments for controlling for the potential selection bias.

The use of panel data requires one to address the initial condition issue in the dynamic demand estimation. By allowing correlation between existing home values and permanent effects, the model allows the initial observed purchase to be potentially endogenous.⁷

2.2 Likelihood function

The second step of the econometric framework models the likelihood of observing household i making J_i home purchases over the periods $\{1, \dots, T_i\}$. Denote $d_i = \{d_{i,1}, \dots, d_{i,T_i}\}$ and $H_i = \{H_{i,t(1)}^1, \dots, H_{i,t(J_i)}^{J_i}\}$, where $t(j)$ indicates the period when the j th purchase occurs. To write down the likelihood function for (d_i, H_i) , I make use of the household-specific permanent effect term μ_i , which enters the timing decision in equation (1) and has a normal distribution with a density function $g(\mu_i)$. Conditional on μ_i , the likelihood of observing (d_i, H_i) for household i is given by:

$$f(d_i, H_i | \mu_i) = f(d_i | H_i, \mu_i) f(H_i | \mu_i), \quad (5)$$

where $f(d_i | H_i, \mu_i)$ is the likelihood function of a T_i -variate conditional probit and $f(H_i | \mu_i)$ is the J_i -variate conditional normal density.

For the ease of notation, define $\mathbf{X}_{it} = \{1, VOL_{it}, VOL_{it} \times HP_{it}, x_{it}, z_{it}\}$ and $\mathbf{Z}_{it} = \{1, VOL_{it}, VOL_{it} \times HP_{it}, x_{it}\}$. For household i , given the J_i transactions it makes, define $\bar{H}_i = \frac{1}{J_i} \sum_{j=1}^{J_i} H_{i,t(j)}^j$ and

$\bar{\mathbf{Z}}_i = \frac{1}{J_i} \sum_{j=1}^{J_i} \mathbf{Z}_{i,t(j)}$. Given the multivariate normality assumption, the second term on the right-hand side of equation (5) can be obtained as:

$$f(H_i|\mu_i) = \exp \left\{ -\frac{1}{2\sigma_\nu^2} \sum_{t=1}^{T_i} d_{it} (H_{it} - \mathbf{Z}_{it}\gamma - \rho_1\mu_i)^2 - \frac{J_i}{2(\sigma_\nu^2 + J_i\sigma_1^2)} (\bar{H}_i - \bar{\mathbf{Z}}_i\gamma - \rho_1\mu_i)^2 \right\} \times A, \quad (6)$$

where:

$$A \equiv (2\pi)^{-\frac{J_i}{2}} \sigma_\nu^{2-\frac{J_i-1}{2}} (\sigma_\nu^2 + J_i\sigma_1^2)^{-\frac{1}{2}} \quad (7)$$

$$\rho_1 = \frac{\sigma_{\alpha\mu}}{\sigma_\mu^2} \quad (8)$$

$$\sigma_1^2 = \sigma_\alpha^2 - \frac{\sigma_{\alpha\mu}}{\sigma_\mu^2}. \quad (9)$$

To obtain an expression for the first term on the right-hand side of equation (5), I rely on the conditional distribution of $\epsilon_{it}|(e_{it}, \mu_i)$. By imposing the multivariate normal distribution, I obtain the conditional expectation of ϵ_{it} as a linear function of e_{it} and μ_i :

$$E(\epsilon_{it}|e_{it}, \mu_i) = d_{it}\rho_2 e_{it} + \rho_3 \bar{e}_i + \mu_i, \quad (10)$$

$$V(\epsilon_{it}|e_{it}, \mu_i) = \sigma_2^2, \quad (11)$$

with:

$$\rho_2 = \frac{\sigma_{\eta\nu}}{\sigma_\nu^2}, \quad (12)$$

$$\rho_3 = -\frac{\sigma_{\eta\nu} \frac{\sigma_\alpha^2}{\sigma_\nu^2}}{\sigma_\nu^2 + J_i\sigma_\alpha^2}, \quad (13)$$

$$\sigma_2^2 = \sigma_\eta^2 - \frac{\sigma_{\eta\nu}^2(\sigma_\nu^2 + \sigma_\alpha^2)}{\sigma_\nu^2(\sigma_\nu^2 + J_i\sigma_\alpha^2)}, \quad (14)$$

and

$$\bar{e}_i = \sum_{j=1}^{J_i} (H_{i,t(j)}^j - \mathbf{Z}_{i,t(j)}\gamma). \quad (15)$$

Adjusting for these correction terms, I write the conditional probability of home purchase timing,

given the observed home purchase values, as:

$$f(d_i|H_i, \mu_i) = \prod_{t=1}^{T_i} \Phi \left[(2d_{it} - 1) \frac{\mathbf{X}_{i,t}\beta + d_{it}\rho_2(H_{it}^j - \mathbf{Z}_{it}\gamma) + \rho_3 \sum_{j=1}^{J_i} (H_{i,t(j)}^j - \mathbf{Z}_{i,t(j)}\gamma) + \mu_i}{\sigma_2} \right], \quad (16)$$

where $\Phi[\cdot]$ is the normal cdf. Substituting equation (6) and (16) into equation (5), and integrating it over the unobserved permanent effect μ_i , I then obtain an unconditional likelihood function for household i :

$$l_i(d_i, H_i) = \int f(d_i|H_i, \mu_i) f(H_i|\mu_i) g(\mu_i) d\mu_i. \quad (17)$$

The final step of the econometric model constructs the likelihood function for all households observed in the sample. Assuming independence across households, the overall log likelihood function for the sample is:

$$\log L = \sum_i \log(l_i(d_i, H_i)). \quad (18)$$

Maximum likelihood estimation is convenient for several reasons. First, when the distributional assumptions hold, the parameters of interest, β_1 , β_2 , γ_1 , and γ_2 , can be consistently and efficiently estimated. Second, all the random-effect parameters are identified. Third, under the null hypothesis of no sample selection bias, $\rho_1 = \rho_2 = 0$. Consequently, the standard Wald test of the significance of additional correcting terms provides evidence of sample selection. Furthermore, by imposing the normality assumption on both equations, one can avoid computational complexities such as multi-dimensional integration.⁸

Although easily implemented, the method above requires strong distributional assumptions.⁹ In addition, the model assumes that all error components are homoskedastic, excluding autocorrelation in idiosyncratic errors in both purchase timing and purchase size equations. Note that the heteroskedasticity in η_{it} does not affect the conditional expectations and therefore can be incorporated by having σ_2^2 vary over time.

3 Data and Variable Construction

3.1 Data

The main dataset used for the estimation of the model is the family questionnaire of the Panel Study of Income Dynamics (PSID) from the Survey Research Center at the University of Michigan, for the period 1980-1997.¹⁰ This is the largest longitudinal U.S. dataset containing both household demographics and housing expenditure information. The longitudinal aspect is particularly useful for identifying the effects of house price risk on housing demand. It allows one to track the observed changes in households' demographic and economic circumstances. In addition, it allows one to control for time-invariant household-specific unobservables that might influence households' expectations and actual decisions.

To construct a sample suitable for estimation, several selection criteria are used. The sample is restricted to existing homeowners; the sample also excludes households that switch back from owning to renting and that stay in a rental house for more than three years.¹¹ And, to preserve enough dynamics, the sample is limited to households that have been observed for more than five years.

There are 2,567 households included in this dataset.¹² All households start with an existing home, and a household's purchase history consists of one or more purchase spells, each spell consisting of one or more years. A purchase spell is complete if it ends with an observed transaction or remains incomplete otherwise. In the sample, over 90% of the households are right-censored in that their last observed purchase spells are incomplete.

Table 1 displays descriptive statistics for the whole sample. For each household in each period, the dependent variables are a dummy variable indicating whether a transaction is made and a continuous variable indicating the purchase size conditional on transacting. The purchase size variable is discussed in detail in Section 3.2.1. As a reflection of lumpy transaction costs, the mean duration between two purchases is 7.53 years. Table 2 illustrates the frequency of home transactions. Among all transactions, 75% are considered "trade-ups" in the sense that the currently purchase house is larger than the previously-owned house in real values; and 69% of moves occur within the same MSA. These

observations suggest a potentially important role for hedging incentives.

The control variables include household demographics, housing user cost, and market conditions. The demographic variables consist of a household head's sex, race, marital status, family income, family size, and number of children. These variables are included to capture a household's housing taste at time t . In addition, the model includes both age and age squared to control for potential nonlinearities in life-cycle home purchase behavior. Following Gyourko and Linneman (1996), educational attainment of the household head is used to indicate the future earnings potential of the family. Furthermore, following Poterba (1984), I construct the user cost term, UC_{it} , to capture prevailing local housing price information (e.g., current price, expected housing return) and market conditions (e.g., mortgage, marginal tax deduction, etc.) that household i faces at time t in the market in which it resides. Appendix 3 explains how the user cost is constructed. In particular, the marginal income rate is computed as the sum of the federal tax rate and the local tax rate; following Flavin and Yamashita (2002), the federal tax rate is assumed to be 28%; and the local tax rates are obtained from income tax data published by Ernst & Young International (the Global Executive, 2000). Finally, I include the MSA fixed effects and year dummies. These controls help capture differences in macro-economic conditions across housing markets that are correlated with local economic activity and also affect the home purchase decision.

As discussed earlier, two sets of instrumental variables are used to control for selection bias. First, the housing stock imbalance measures the difference between the desired housing stock and the existing housing stock. Since the desired housing stock is not observed for "stayers" (i.e., those who decide to not to transact in year t) and the existing housing stock is potentially endogenous, I proxy the housing stock imbalance by changes in household demographics between the current year and the year in which the previous transaction occurred. In particular, I consider the following variables: changes in family income, changes in the head's marital status, changes in the head's working status, and increases and decreases in the number of children.

A second set of instrumental variables relates to transaction costs. Following the housing literature, the paper uses 10% of the existing home value to proxy monetary transaction costs. This is meant

to capture transaction taxes and broker fees involved in selling a house. Additionally, a waiting time variable and its square value are included to capture possible psychic costs in home transactions. Intuitively, a longer stay in the current house may be associated with a stronger connection with local neighborhood facilities and hence may make it more costly for a household to move.

The main dataset of the PSID has been supplemented by data from two additional sources. First, while the public-use version of the PSID provides the household's residence information only up to the state level, the restricted-access geocode files of the PSID are used to identify the residence of each household at the metropolitan area level. Using a less aggregated market helps reduce the measurement error in constructing a household's forecast over future economic circumstances.

As a second source, the paper uses the Conventional Mortgage Home Price Index (CMHPI), published by Freddie Mac, to construct housing return and risk variables. The CMHPI provides quarterly, nominal weighted repeat-sales house price indices on single-family homes, beginning in 1974. These data become available for 148 metropolitan areas, beginning in 1980. The CMHPI data enable one to estimate a time-series of housing returns and the associated risks for each local housing market.

3.2 Variable construction

3.2.1 Purchase size

For the demand analysis to be meaningful, it is critical to have a measure of house size, H_{it} , which meets the following two requirements. First, in theory, H_{it} is an argument of the utility function, and should therefore reflect some measure of the physical quantity of housing rather than expenditure. Second, in practice, housing is a differentiated product made up of a set of attributes that include structural characteristics and local amenities. An ideal measure of H_{it} should permit comparison of the housing stock from one period to the next period and from one market to another market.

Absent a direct measure of housing quantity, I impute housing quantity by deflating house gross value with a quality-adjusted house price index. The resulting housing quantity measure is not only consistent with theory but also comparable over time and across markets. As an empirical check, I also

consider three alternative measures of home purchase size – gross house value, number of rooms, and imputed square footage. Each provides a useful complementary test for the main empirical analysis. This section discusses how each of these four measures is constructed.

I start by describing how to impute housing quantity. In principle, one could begin with PSID data on the value of the house and divide the value by the unit price of housing stock. Since there are no data on the unit price of housing stock, I use the constant-quality house price index instead. Existing repeat-sales house price indices control for changes in house quality over time, but not across markets. For example, all OFHEO house price indices are set to 100 in the first quarter of 1995 and hence cannot be directly compared across MSAs. To remedy this, I use the MSA-level hedonic house value estimates based on 1990 Census data (Malpezzi, Chun, and Green (1998)) and apply a cross-sectional hedonic adjustment to the OFHEO repeat sales house price indices. This permits me to obtain a hedonic-adjusted house price index that is comparable both over time and across markets. I then impute the home purchase size by deflating the gross house value reported by the PSID with the hedonic-adjusted house price index. Appendix 4 describes in detail how the hedonic-adjusted house price index is constructed. The resulting housing quantity variable is consistent with its theoretical content. To the best of my knowledge, this is the first quality-adjusted housing demand measure that is comparable both over time and across markets.

Despite its conceptual appeal, the imputed housing quantity variable may omit variation in house prices within an MSA for a given year. For this reason, I consider three alternative measures of housing demand as robustness checks: gross house value, imputed square footage, and number of rooms. None of these three measures is perfect. The gross house value reports the total expenditure on housing rather than quantity, while the imputed square footage, and number of rooms are potentially contaminated by unmeasured variation in other housing attributes. Nevertheless, using these three measures avoids the possible aggregation bias that arises from using the MSA-level price index, and therefore provides useful checks for testing the predictions of the model.

Among all measures of housing demand, house value is used most often in empirical housing demand studies (e.g., Goodwin (1986), Bajari, Chan, Krueger, and Miller (2008)). This is because house value

conveniently summarizes information on housing characteristics, including physical size. Following the convention in the literature, I use the gross house value, self-reported by PSID homeowners, and then convert it into real terms using the CPI-U deflator as an alternative measure of home purchase size. It is reasonable to question how precise homeowners' self-reported house values are. Skinner (1994) compared self-assessed house values in the PSID to objective measures from the Department of Commerce and found the two series to be quite close in mapping housing price changes in the 1970's and 1980's. Despite the accuracy and convenience of using the self-reported house value, it is worth noting that this measure is proportional to the quantity of housing stock only if all households face the same house price schedule. Such a condition is unlikely to hold given that housing is highly differentiated and location-specific. When a transaction occurs, changes in housing value may reflect changes in housing structure, as well as changes in locations.

For additional robustness checks, I consider two physical measures of housing quantity. The simplest quantity measure of housing is the number of rooms in a house. This variable is appealing for two reasons. First, it is self-reported by the PSID homeowner and does not require any additional imputation. Second, it provides an easy interpretation when measuring the physical amount of housing. However, the number of rooms provides only a rough approximation to the size of home purchase as it does not take into account variation along many other dimensions, such as the size of rooms.

Finally, I consider another physical measure of housing quantity: the square footage of a house. The square footage data are not reported by the PSID, but can be imputed using an approach described by Flavin and Nakagawa (2008). They used data from the American Housing Survey (AHS) to estimate a model of square footage as a function of the number of rooms and other housing variables common to both the AHS and the PSID. Using their estimates from the AHS sample, I impute the square footage for each household observation in the PSID sample. This provides an alternative measure of housing quantity for the demand analysis. While square footage provides a better measure of purchase size than the number of rooms, it omits variation in locations and building materials.

Table 3 presents summary statistics for four different measures of housing demand in the PSID sample. The mean house value is about \$77,577 (in 1982 dollars). After being deflated by the hedonic-

adjusted house price indices, the mean house size is about 1,395. Table 4 presents the Spearman correlation among all four measures. As one might expect, all correlation coefficients are significantly positive. If we take the hedonic-adjusted house size as a proxy for the “true” quantity of housing, the magnitude of these correlation coefficients reveals the degree of consistency between alternative measures of housing quantity and quality-adjusted housing quantity. Not surprisingly, none of the three alternative measures is perfectly consistent with the “true” quantity, as their correlation coefficients are less than unity. It is perhaps more surprising that the value measure is much more consistent with the “true” quantity than the physical measures (the imputed square footage and the number of rooms). This is probably because the former contains more information about housing structures, as well as the services that flow from local amenities.

3.2.2 Hedging propensity

The key to testing the hedging effect of price risk on housing demand lies in the ability to measure hedging propensity. As shown in Section 1, a homeowner’s hedging propensity should capture two ingredients: (1) the probability of trading up to a bigger house in the future; and (2) the positive correlation between the price of the currently desired house and the price of the future desired house. To proxy hedging propensity, this paper assumes that, based on the information in period t , household i predicts its probability of trading up to a bigger house within the same market in the next five years.¹³ Intuitively, the more likely a household plans to move within the same market, the more likely the prices of its current and future houses will be positively correlated, and the stronger its hedging incentive is. Thus, the predicted hedging propensity is consistent with the theoretical characterization of hedging incentives for homeowners.¹⁴ However, one disadvantage of relying on geographical mobility is that it implicitly assumes away variation in hedging incentives for households that plan to move across markets. As shown by Sinai and Souleles (2008), many housing markets are positively correlated with each other, thus strong hedging incentives could be present even among households that move across these markets. The hedging propensity constructed above does not account for hedging incentives for cross-market movers. In this sense, the estimated hedging effect in this paper should be treated as providing a lower bound for the true hedging effect.

Figure 2 describes the decision tree for modeling the hedging propensity. Conditional on purchasing a home now, the first level of the decision tree predicts whether to trade up to a bigger house in the next five years. This is represented by a dummy variable, $U_{i,t+5}$, which equals one if $d_{i,t'} = 1$ and $H_{i,t'} > H_{it}$, where $t < t' \leq t + 5$. Conditional on trading up, the second level of the decision tree predicts whether the future move occurs within the same market. This is represented by another dummy variable, $W_{i,t+5}$. To integrate these two decisions, I estimate a sequential logit model. In particular, a household's hedging propensity is measured by the conditional probability of moving to a bigger house within the same MSA, which is given by:

$$Pr_{it}(U_{i,t+5} = 1, W_{i,t+5} = 1 | d_{it} = 1) = Pr_{it}(W_{i,t+5} = 1 | U_{i,t+5} = 1, d_{it} = 1) \times Pr_{it}(U_{i,t+5} = 1 | d_{it} = 1), \quad (19)$$

where:

$$Pr_{it}(U_{i,t+5} = 1 | d_{it} = 1) = \frac{e^{w_{1,it}\lambda_1 + \alpha_i}}{1 + e^{w_{1,it}\lambda_1 + \alpha_i}}. \quad (20)$$

$$Pr_{it}(W_{i,t+5} = 1 | U_{i,t+5} = 1, d_{it} = 1) = \frac{e^{w_{2,it}\lambda_2}}{1 + e^{w_{2,it}\lambda_2}}. \quad (21)$$

The subscript it refers to the fact that the expectation is conditional on household i 's information in year t . Note that the definition of trading-up ($U_{i,t+5} = 1$) depends on the size of the future house relative to the size of the current house, where the house size is constructed in four different ways in Section 3.2.1.

The vector $w_{1,it}$ consists of household i 's characteristics – age, age squared, head's education, education squared, job tenure, job tenure squared, marital status, family size, family size squared, and family income – in year t . In addition, using panel data allows one to include the household fixed effect, α_i , which is known to households but not observed to econometricians.

The vector $w_{2,it}$ contains variables in $w_{1,it}$ and a set of instrument variables, such as whether the head's occupation is licensed, along with age composition dummies for children. These variables are included to control for the potential endogeneity associated with using the hedging propensity in the main demand model. Section 4.1 provides justification for using these instruments.

3.2.3 House price risk

Another challenge in this empirical analysis is that an appropriate empirical measure of house price risk is not obvious. Previous housing demand studies have proxied price risk with either location indicators (Ioannides and Rosenthal (1994), Green and Vandell (1999)) or with variability in past house price history (Sinai and Souleles (2005)).¹⁵ The location proxies are sensitive to measurement error and do not capture time variation in local market conditions, while the variability approach often requires the truncation of the initial periods of more than five years, which represents a substantial loss of information. More importantly, as with other investments, housing decisions depend on a household's expectation of future economic circumstances at the time when these decisions are made, rather than on the observed past history of house prices alone. This implies that the measure of price risk in the housing demand model should vary over time and do so in a way that is at least partly predictable.

This paper assumes that, at the beginning of each period, households make forecasts about future housing returns and volatility based on past house price history and local market conditions. In particular, the determinants of house price appreciation are divided into two parts: one that accounts for the influence of the local market conditions and another that accounts for the adjustment dynamics. The market condition part includes lagged changes in income per capita and lagged changes in the unemployment rates in local markets. I obtain local income and unemployment rates from the Bureau of Economic Analysis Regional Accounts data. Given the sample size in each market, I am constrained from including more market factors in the regression. However, by allowing model parameters to vary across markets, the model explicitly controls for market fixed effects. The price adjustment part is modeled by an AR(1)-GARCH(1,1) process. As shown by Bollerslev (1986), the virtue of the GARCH model is that a small number of terms appear to perform as well as or better than an ARCH model with many lags. This justifies the use of one lag in our empirical analysis. In the simplest case where all past variances receive equal weights in predicting future variance, the GARCH measure of price risk is equivalent to the average standard deviation measure.

Specifically, I express the quarterly housing return at time t in market m , $RET_{m,t}$, as a function of lagged income growth, $INC_{m,t-1}$, lagged unemployment rate changes, $UE_{m,t-1}$, lagged housing return,

$RET_{m,t-1}$, and a return shock, $u_{m,t}$. The conditional variance, $VOL_{m,t}$, is modeled as a function of previous return shock, $u_{m,t-1}$, and previous conditional variance, $VOL_{m,t-1}$. For each metropolitan area, I estimate the following AR(1)-GARCH(1,1) process separately:

$$RET_{m,t} = a_{m,0} + a_{m,1}RET_{m,t-1} + a_{m,2}INC_{m,t-1} + a_{m,3}UE_{m,t-1} + u_{m,t} \quad (22)$$

$$VOL_{m,t} = b_{m,0} + b_{m,1}u_{m,t-1} + b_{m,2}VOL_{m,t-1} \quad (23)$$

where $u_{m,t} \sim N(0, VOL_{m,t})$. Note that in each period, the return shocks, captured by $u_{m,t}$, cause realized returns to deviate from expectations. For example, $u_{m,t}$ could be negative if there is a sudden increase in the local unemployment rate. This new information is used to predict next period's return and conditional volatility.

Table 5 reports cross-sectional averages of the estimated coefficients. The results have the expected signs. For example, the cross-sectional averages of $a_{m,1}$ s are statistically positive, suggesting that housing markets are inefficient in the short run (Case and Shiller (1989)). In addition, consistent with what one might expect, house price appreciation is positively correlated with local income growth but negatively correlated with the local unemployment rate. Finally, the GARCH coefficients are significant in most markets, suggesting the importance of modeling the adjustment process in predicting future housing prices.

The estimates generate conditional forecasts of quarterly housing return and volatility. Using the time aggregation procedures explained in Mwang and Quigley (2006), I generate forecasts of the five-year housing return and associated risk.¹⁶ The results are then mapped to the households in the PSID sample at the MSA level for each year. The predicted five-year conditional variance is used as the proxy for house price risk VOL_{it} , while the predicted five-year conditional return is used as the proxy for housing return that enters the user cost term.

4 Identification Strategies

The key variables of interest, such as hedging propensity, price risk, and transaction costs, are likely to be endogenous when carrying out housing demand estimation. To obtain consistent estimates, this section presents identification strategies to address each of these potential concerns.

4.1 Endogeneity of hedging propensity

A central challenge in developing a compelling test of the hedging hypothesis is that the hedging propensity is likely to be endogenous to a household's response to price risk. By adopting a correlated random effects specification, the demand model controls for the time-invariant unobservables that could affect both the hedging propensity and housing demand. However, this approach cannot address the endogeneity concerns associated with unobserved *time-varying* shocks. For example, when a household experiences a local economic boom, the household is likely to plan to move up and hence, has stronger hedging incentives. The same economic boom could also make the household less risk-averse and push it towards a bigger home purchase in a volatile market. In this case, one would find a positive coefficient on the interaction term between the hedging propensity and price risk even in the absence of any hedging effect.

To address this particular endogeneity concern, I use the following variables to instrument for hedging propensity: whether the head's occupation is licensed, as well as composition dummies for the ages of children. The discussion below suggests that these instrumental variables proxy for the monetary and psychic costs of inter-market movements but do not affect housing demand directly, and therefore, create exogenous variation in hedging propensity.

First, occupational licensing directly affects approximately 12% of U.S. workers. Some occupations require obtaining local licenses or developing a local reputation, which could significantly restrict licensees' geographical mobility. For example, a lawyer's family is more likely to move within a state than a university professor's family. To control for this effect, I have created a dummy variable 'licensed occupation,' which indicates whether the household head's occupation is licensed in its local market.¹⁷

Second, among households with children, geographical mobility is lower for families with school-aged children, since the moving cost associated with disrupting their education is high. To capture such variation in hedging incentives, I have created a set of binary variables representing the presence of children of specific ages. The age categories are 0-6 (preschool age children), 7-12 (elementary school age children), and 13-18 (teenagers).

Of course, these variables are valid instruments only if they do not have direct effects on housing demand conditional on other covariates in the main specification. Since the demand regression has explicitly controlled for the number of children (both in levels and in changes), it is unlikely that the age composition of children would have any effect on housing demand other than through the hedging channel. On the other hand, it is less obvious whether having a licensed occupation has any direct effect on housing demand. Section 6.2 tests whether this exclusion restriction holds.

4.2 Endogeneity of price risk

Another key variable of interest is house price risk. Following the existing literature, this paper takes the movements in house prices as given. The standard reversal causality issue is less of a concern in this paper than in other settings. First, the demand model requires expected housing returns and risks rather than actual house price changes. The GARCH model predicts housing returns and risks based on the *past* history of local house prices and market conditions, which are unlikely to be affected by *current* household level demand shocks. Second, by adopting the correlated random effect specification, the model has controlled for the possible bias that is caused by market-specific unobserved conditions.

However, expected housing returns and associated risks are likely to be tied to changes in local labor market conditions (Spiegel (2001)), which along with housing demand, may be jointly determined by some *time-varying* local market unobservables. Without further modeling, this could potentially bias the estimated price risk effects. For example, in the second half of the 1980s, the San Jose, California economy experienced an unprecedented expansion that was driven by the Silicon Valley IT industry boom. As a result of better employment opportunities and higher income enjoyed by residents in this expanding metropolitan market, both local house prices and housing demand increased. In this case,

one might find a significant correlation between price risk and housing demand, even in the absence of a causal effect of risk on demand.

To control for this type of endogeneity issue, I present a test of house price endogeneity using a control function approach (Blundell and Powell (2004)). I use two variables as instruments: lagged changes in construction worker income and lagged changes in the number of construction workers. These two variables are correlated with households' expectations over future house price movements through the supply side but are uncorrelated with their home purchase decisions (Somerville (1999)), making them suitable instruments. A description of the endogeneity test is given in Appendix 5.

4.3 Endogeneity of transaction costs

I use 10% of the existing home value as a proxy for transaction costs. One concern is that a household's current housing demand may be correlated with its previous housing demand through unobserved tastes. Since the correlated random effects model has allowed for the correlation between unobserved permanent tastes and transaction cost variables, the identification assumption becomes the following: a household's existing home value is independent of the realization of its current demand shock.

This assumption is justified by the dynamic structure of home purchases. The size of the existing home is inherited from the previous home purchase, which occurred, on average, seven years ago.¹⁸ In the simplest case where there is no serial correlation in transitory shocks, it is easy to argue that the current shock would not affect existing home size. Even in the case where the transitory shocks follow an AR(1) process, one could reasonably expect that the serial correlation between the current shock and the shock seven years ago would gradually die out as time passes.

5 Empirical Findings

5.1 First-stage results: hedging propensity estimates

Table 6 reports results from estimating the hedging propensity in the first stage. I begin by discussing the results in Specification I, where house size is measured by the imputed hedonic-adjusted housing quantity.

The first stage of the sequential logit model estimates the probability of a household trading up to a bigger house within five years. Being younger and being more educated are both statistically significant and positive predictors of trading-up activity with decreasing marginal effects. In addition, households with higher income are more stable and less likely to move again. Job indicators, such as job tenure, have no noticeable effect on households' future housing consumption plans.

Conditional on trading up, the second stage of the sequential logit model estimates the probability of moving within the same market. Several findings emerge. First, households whose heads engage in licensed occupations are more likely to move within the same market – an effect that decreases with the head's job tenure. The findings are consistent with studies that have highlighted the importance of occupation differentials in inter-market mobility (Pashigian (1979) and Kleiner, Gay, and Greene (1982)).

Second, everything else being equal, households without children have the highest probability of migrating across markets, followed by households with preschool-aged children, households with teenagers, and finally, households with elementary-school-aged children. Intuitively, among households with children, those with only preschool-aged children have more freedom to move across markets because their children are not yet involved in particular school districts. The opposite is true for households with school-aged children.

Finally, turning to the remaining coefficients, I find that older and more educated households are more likely to move across markets. This is probably because they face more job opportunities or because they are more responsive to regional wage differentials. In addition, households with two income earners are less likely to move across markets. This can be justified by the high costs of

simultaneously searching for two jobs in a different market.

The first-stage estimates of the hedging propensity are of interest in their own right, since they demonstrate significant heterogeneity in hedging behavior among different households and across different stages of the life-cycle. For example, among families that live in the same market and have similar family characteristics, a 35-year-old lawyer’s family would have stronger incentives to hedge than a 35-year-old professor’s family due to differences in their geographical mobility. Among lawyers’ families, younger families are more likely to hedge than older families, given that the former are more likely to move up the housing ladder. The heterogeneities in hedging propensities provide a key source of variation that helps identify the price risk effects in the second-stage demand estimation.

As an empirical check, Specifications II-IV in Table 6 repeat the estimation in Specification I for alternative measures of house size, which is a key variable in defining whether a household trades up in the next five years. The results are qualitatively consistent with the estimates in Specification I. In particular, the life-cycle variables, such as age and education, continue to have strong positive effects on the probability of trading-up in Specification II, where house size is measured by the gross house value. The corresponding estimates are statistically and economically weaker in Specifications III and IV, where house size is measured using the imputed square footage and the number of rooms, respectively. This is not surprising given the previous finding that house value is more positively correlated with quality-adjusted housing quantity than single-attribute measures of house size.

For each specification reported in Table 6, I first construct a continuous measure of the ‘hedging propensity score’ (HP_{it}) by computing the predicted probability of trading up to a bigger house within the same market within five years. The measure is a probabilistic variable between 0 and 1. I then generate two dummy variables that separate households equally into three groups: low hedging propensity, medium hedging propensity, and high hedging propensity. I assign a value of 0 or 1 for the “medium hedging propensity” (HP_{it}^m) and “high hedging propensity” (HP_{it}^h) to each household in each year of the sample.

Table 7 reports descriptive statistics for various household characteristics, broken down by the hedging propensity dummies computed using the estimates in Specification I from Table 6. The statis-

tics reveal significant differences among households with low, medium, and high hedging propensities. For example, younger households with a smaller family have stronger hedging incentives. In addition, households with medium levels of education and income have the strongest hedging incentives. Intuitively, this may indicate that low-income families are financially constrained from moving up and that high-income families are more stable and less likely to move up again. The general pattern is consistent with an inverted-U shape in households' hedging behavior across the life-cycle.

5.2 Second-stage results: housing demand estimates

The main estimates from the dynamic housing demand model are presented in Table 8. Conditional on transacting, purchase size is measured by the imputed hedonic-adjusted housing quantity. Specification I reports demand estimates based on a continuous measure of hedging propensity. In both purchase timing and purchase size equations, housing price volatility (VOL_{it}) has significantly negative coefficients. This indicates that, absent hedging incentives, an increase in price risk not only delays home purchases, but also reduces the size of a home purchase if transacting. Moreover, the positive coefficient on the interaction term ($VOL_{it} \times HP_{it}$) indicates that, as the level of hedging propensity increases, the negative effects of price risk on housing demand are mitigated or even reversed. These results are consistent with the predictions from the theory.

Specification II in Table 8 reports demand estimates using discrete measures of hedging propensity. This is meant to capture possible nonlinear effects of the hedging propensity. The results are generally consistent across specifications. To see whether households with different levels of hedging propensity respond to price risk differently, I test $\beta_{HP^mVOL} = 0$, $\beta_{HP^hVOL} = 0$ and $\beta_{HP^m} = \beta_{HP^h}$ jointly. The Wald test statistic shows that the hypothesis is rejected at the 5% level. This implies that households' risk responses differ across households in a way that depends on their hedging propensities. To make the coefficients easy to interpret, I measure the housing return risk as the square root of the forecasted conditional variance, i.e., $\sigma_{it} = \sqrt{VOL_{it}}$. By doing so, the expected house price risk is defined in the same units as the expected housing return. Specifically, a one percentage-point increase in the five-year housing price risk is interpreted as shifting the house price risk one percentage-point from the

sample mean of σ_{it} . To quantify the importance of heterogeneity, Table 9 reports the marginal effects of house price risk for three groups of households, with low, medium, and high hedging propensities, respectively. On the timing side, a one percentage-point increase in the price risk *decreases* the transaction probability by 0.68% for households with low hedging propensities, and by 0.34% for those with medium hedging propensities; but it *increases* the transaction probability by 0.45% for those with high hedging propensities. Conditional on transacting, a one percentage-point increase in the price risk *decreases* the purchase size by 1.42% for those with low hedging propensities and by 0.11% for those with medium hedging propensities. For those with high hedging propensities, it *increases* the purchase size by 1.06%. These estimates provide strong support for the hypotheses laid out in Section 1.

Since price risk is defined as the standard deviation of the expected return, it is natural to provide a comparison of the effects of the expected housing return on existing homeowners' housing demand. As described in Appendix III, the expected housing return is included as part of the housing user cost. Column 2 in Table 9 reports the marginal effects of the expected housing return based on the estimates from Specification I in Table 8. A one percentage-point increase in the expected return increases the conditional purchase size by 1.75% and increases the transaction probability by 0.93%. Thus, for households with high hedging propensities, their hedging demand is fairly comparable to the investment demand motivated purely by expected returns.

Turning to other variables in the demand estimation, I find that being married, being highly educated, and having a higher income increase both the probability of transacting and the size of purchase conditional on transacting. Households whose heads are older and who have worked relatively longer tend to make less frequent transactions but bigger purchases. The estimates are consistent with a life-cycle model of home purchase behavior. In addition, in both purchase timing and size equations, the coefficients on the user cost are significantly negative, indicating a downward-sloping demand relationship.

Transaction cost variables also play an important role. For example, an increase in the financial transaction costs, proxied by 10% of existing home value, leads to a decrease in the transaction proba-

bility. Waiting time variables are included to capture the unobserved factors that increase households' transaction costs over time. The results show strong negative duration dependence with decreasing marginal effects over time. This confirms the intuition that, everything else equal, a longer stay in the current home would lead to a stronger connection with the local community, which makes it more costly for a household to move.

Turning to proxies for housing stock imbalance, both changes in marital status and changes in working status increase the probability of transaction. In contrast, changes in family income do not have much effect. One interpretation is that changes serve as good proxies for housing stock imbalance only when they are unpredictable. This is because predictable changes have been taken into account by forward-looking households in their previous home purchase decisions, and hence do not reflect the current housing stock imbalance. Since income changes are mostly predictable according to the permanent income hypothesis, they should not affect current purchase decisions.¹⁹ I also find that changes in the number of children have asymmetric effects. For example, increases in the number of children have a positive and statistically significant effect on the transaction probability, but decreases in the number of children do not appear to affect transaction probability. While the former is clearly an indicator of the housing stock imbalance effect, the latter probably indicates that parents are reluctant to downsize the house even after their children move out.

Finally, as shown in the bottom of Table 8, there is substantial evidence of unobserved permanent effects. The positive coefficient on the permanent level selection term implies that those who are likely to make home purchases also prefer to purchase a bigger home if transacting. The significance of household permanent effects and of the associated selection effects suggests that panel data are indispensable in the empirical analysis of existing homeowners' housing demand. Furthermore, the correlated random effect model specifies the permanent effects as linear functions of the average lagged and forecast variables over the observed periods (\overline{TC}_{it} , \overline{HP}_{it}), where the average is taken over all the observed periods for each household. The estimates for average transaction costs and average hedging propensities are both statistically significant, suggesting that a model that does not account for the correlated random effect structure could lead to biased results.

5.3 Sensitivity to Specification

To probe our results further, I fit a series of specifications based on equations (1) and (2) that use alternative measures of housing size and additional instruments from the housing market supply side. Table 10 presents estimation results. The top panel reports estimates from a joint housing demand model where a continuous measure of hedging propensity is used. The second panel reports estimates from the same model where discrete measures of hedging propensity are used.

The results in Specification I of Table 10 replicate the baseline estimation in Table 8, where home purchase size is imputed by deflating the gross house value using the hedonic-adjusted house price indices. Specifications II-IV repeat the estimation in Specification 1 based on alternative measures of home purchase size. In Specification II, where the deflated house value is used to measure house size, the risk estimates remain statistically and economically significant. This is consistent with the descriptive evidence that the two measures of house size are highly positively correlated. Although the house value itself does not control for spatial variation in house prices, the MSA-fixed effects in the demand estimation help pick up the time-invariant differences in local amenities across metropolitan areas.

In Specification III-IV, where the purchase size is measured by the imputed square footage and the number of rooms respectively, the risk estimates become much weaker in magnitude and in significance. This is not surprising given that these single-attribute physical measures of housing quantity do not control for unmeasured variation in other attributes and local amenities. Despite the difference in magnitude, the signs of the estimates in Specifications II-IV are consistent with those in Specification 1, suggesting that the main results in this paper are qualitatively robust to the different measures of housing size.

The estimates shown in Specifications I-IV are identified assuming that a household's house price forecast is independent of its current demand shocks. As discussed above, it is possible that local housing market conditions are correlated with variation in the local labor market. If the latter is the driving force for the demand shocks, the estimates in Specifications I-IV would be biased. Specification V repeats the estimation in Specification I but allows a linear dependence between price variables and

the household permanent effect term using the correlated random effect approach described above. The results again are consistent with the benchmark results in Specification I, providing additional support for the predictions from the theory.

To further control for the endogeneity associated with price risk, the paper uses an instrumental variable strategy where the instruments are chosen from the supply side of the housing market. Specification VI in Table 10 presents the estimates of risk effects using the following instruments: lagged changes in construction worker income and lagged changes in the number of construction workers, both of which change over time and across markets. Controlling for these instruments does not show any substantial effect on the estimates of house price risk terms. This is probably because that the correlated random effect specification has already largely controlled for the bias that is caused by market-specific unobservables.

So far, I have identified the heterogeneity in households' risk responses by relying on variation in their hedging propensities. An alternative source of heterogeneity in households' responses to price risk can be explored by looking at the background risk in home purchase decisions.²⁰ One way to explore this is to examine whether the price risk effects vary with households' investment in housing futures. However, this is not feasible in our empirical analysis, as the housing options and futures were unavailable during the sample period of 1980-1997. Another source of background risk is the entrepreneurial risk faced by private business owners (Heaton and Lucas (2000)). Since households that own private business may be better able to diversify their housing investment risk, one might think that the strength of the financial and hedging incentives varies with the business ownership. To test the differential effect of price risk by business ownership, I repeat the baseline estimation in Table 8 by including a dummy variable indicating whether the head of a household owns a private business, its interaction with the price risk, and a triple-interaction term between the business ownership dummy, price risk and hedging propensity.

The estimates are reported in Table 11. The results suggest that owning a private business increases the size of purchase conditional on transacting. The coefficients on the interaction terms are insignificant, indicating that being a business owner has no direct effect on hedging demand. Recall that

Table 6 reports a significant role for the self-employment dummy in the first-stage regression. Since a business owner is likely to be self-employed, I infer that owning a business could increase hedging effects indirectly by raising the predicted hedging propensity.

6 Robustness Checks

To provide further evidence on the strengths and weaknesses of the model, I now present estimates from several additional specifications to examine the following assumptions that have been made in the paper: (i) that the hedging propensity is not correlated with risk aversion, and (ii) that the occupation instrument does not directly affect housing demand.

6.1 Risk aversion

The empirical evidence supports the theory of hedging demand only if other reasons for such a link are ruled out. Kan (2003) uses the survey questions on risk attitude in the 1996 PSID supplementary file and finds that risk aversion enlarges the risk effects by a small but non-negligible magnitude. This raises the question: Can the omitted risk aversion factor be responsible for the findings of the hedging effect? Unfortunately, the information on risk attitudes is available for only one year in the PSID, and hence cannot be explicitly controlled for in our estimation.

To address this question, I separately discuss two types of difference in a household's risk aversion that may be responsible for the hedging effect: time-invariant differences across households and time-varying differences for the same household. The time-invariant differences are less of a concern in this paper given the correlated random effect specification. A more serious concern is due to possible variation in risk aversion that changes with age. For example, compared to older households, young households tend to be disproportionately less risk-averse and more hedging-motivated. The unobserved age-varying changes in risk aversion confound identification, since the effect of an increase in a household's hedging incentive on its demand coincides with the effect of a reduction in risk aversion. In other words, the positive hedging effect that is uncovered for young households could in theory be

the by-product of a lower risk aversion effect.

To assess whether this is a serious problem, I propose a test to demonstrate that age-varying risk aversion is not the driving force behind our empirical findings. In particular, I first remove the age variable from the first-stage predicted hedging propensity and then re-estimate our joint demand model with this hedging propensity. The hope is that this removes, or at least attenuates, the mechanism through which age-varying risk aversion affects hedging incentives, but leaves intact other factors that can predict hedging incentives. If age-varying risk aversion was the driving force for the previous results regarding hedging effects, then this specification should not generate much of a hedging effect. The results are reported in Specification 7 of Table 12. The point estimates of the hedging effect fall but are still significantly positive, suggesting that the main results are robust even after accounting for the possible risk-aversion effect.

6.2 Validity of instrumental variables

The ideal instruments for hedging propensity induce exogenous variation in households' future housing consumption plans but are orthogonal to unobserved heterogeneity in current housing demand while holding other controls fixed. Table 6 presents evidence that the licensed occupation instrument is a strong predictor of the hedging propensity. However, one may question that this variable also enters housing demand decisions directly. The concern is that human capital acquisition would matter for occupational choice and be correlated with a household's response to price risk, say, through risk attitudes or housing taste.

To test whether the exclusion restriction holds for the occupation instrument, I examine a control group where the first-stage relation between the occupation license indicator and hedging propensity is shut down for intuitive reasons. Specifically, I use the sample of households whose heads are retired as the control group. For retired people, the occupation license indicator no longer ties them to a certain location. Hence, there is no reason to expect a first-stage relationship between occupation and hedging propensity. Repeating the estimation in Table 6 for retired households, I find that this prediction holds true in the data: the occupation license indicator has no significant effect in predicting the hedging

propensity for retired households. Thus, the first-stage relationship is effectively shut down for this control group. The lack of a first-stage relationship means that any association between the occupation and housing demand in this group would constitute evidence of a direct link between the occupation license indicator and housing demand, violating the exclusion restriction. However, repeating the estimation in Table 8 with a licensing indicator, I find that the results reveals no such association. These results support the exogeneity of the occupation instrument.²¹

7 Conclusion

This paper estimates a dynamic infrequent home purchase model for existing homeowners. The empirical analysis extends the prior work that studies the effects of house price risk on housing demand in several aspects. First, this paper presents a simple and intuitive way to measure the hedging propensity of a household, drawing attention to two important and related considerations: geographical mobility and the tendency to trade up. Second, it also addresses the endogeneity problems associated with including the hedging propensity in the demand estimation. By doing so, the paper has demonstrated that the estimated hedging effect cannot be easily explained away by unobserved factors that affect both households' current home purchase decisions and their future housing consumption plans. Third, because the demand specification captures the fundamental features of a housing market – lumpy transaction costs and forward-looking behavior – the model serves not only as a test of homeowners' hedging behavior, but also provides evidence regarding the general importance of illiquidity and dynamics in home purchase behavior.

The paper finds two distinct effects of price risk on housing demand: in response to financial incentives, households reduce current housing demand to avoid future financial risk; and in response to hedging incentives, households take a bigger housing position to offset potentially large housing costs in the future. In a volatile housing market, whether a home purchase is perceived as risky depends on whether the current purchase serves as a hedge against future housing consumption risk. The evidence of self-insurance behavior is important as many households anticipate stepping up the property ladder, but lack financial instruments to manage house price risk. Even if markets were complete, mitigating

price risk through the functioning of financial markets could be very costly, as noted by Englund, Hwang, and Quigley (2002). The findings suggest that, while families in incomplete housing markets are not able to access formal financial instruments, they do rely on private self-hedging mechanisms to reduce house price risk.

Appendices

Appendix 1: Correlated Random Effect Model

For ease of notation, I define $\mathbf{X}_{it} = \{1, VOL_{it}, VOL_{it} \times HP_{it}, x_{it}, z_{it}\}$ and $\mathbf{Z}_{it} = \{1, VOL_{it}, VOL_{it} \times HP_{it}, x_{it}\}$. Denote $\mathbf{X}_{it} = (\mathbf{X1}_{it}, \mathbf{X2}_{it})$ and $\mathbf{Z}_{it} = (\mathbf{Z1}_{it}, \mathbf{Z2}_{it})$, where $\mathbf{X1}_{it}$ and $\mathbf{Z1}_{it}$ are time-varying factors, including hedging propensity, transaction costs, price risk, and so on; and $\mathbf{X2}_{it}$ and $\mathbf{Z2}_{it}$ are time invariant factors, such as race.

The basic idea of correlated random effects is to approximate the conditional expectations $E(\mu_i|\mathbf{Z}_i)$ and $E(\alpha_i|\mathbf{X}_i)$ by linear functions of time-varying regressors:

$$\mu_i = \overline{\mathbf{X1}}_i' b_1 + u_i \quad (\text{A-1})$$

$$\alpha_i = \overline{\mathbf{Z1}}_i' b_2 + a_i \quad (\text{A-2})$$

where $\overline{\mathbf{X1}}_i = \frac{1}{T} \sum_{t=1}^T \mathbf{X1}_{it}$ and $\overline{\mathbf{Z1}}_i = \frac{1}{T} \sum_{t=1}^T \mathbf{Z1}_{it}$. $u_i \sim N(0, \sigma_u^2)$ and $a_i \sim N(0, \sigma_a^2)$ are residual permanent effects. $\overline{\mathbf{X1}}_i$ and $\overline{\mathbf{Z1}}_i$ can be generalized so that they include all the time-varying regressors in \mathbf{X}_i and \mathbf{Z}_i . In the baseline specification, I simplify the estimation by letting $\overline{\mathbf{X1}}_i' = \{\overline{VOL}_i, \overline{VOL} \times \overline{HP}_i, \overline{TC}_i\}$ and $\overline{\mathbf{Z1}}_i' = \{\overline{VOL}_i, \overline{VOL} \times \overline{HP}_i\}$. Clearly, b_1 and b_2 will be equal to zero (and $\sigma_u^2 = \sigma_a^2$ and $\sigma_\alpha^2 = \sigma_a^2$) if and only if the included time-varying regressors are uncorrelated with the permanent effects.

Summarizing equations (3),(4), (A-1), and (A-2), I impose the following distributional assumption:

$$\begin{pmatrix} \epsilon_i \\ e_i \end{pmatrix} = \begin{pmatrix} \mu_i \iota + \eta_i \\ \alpha_i \iota + \nu_i \end{pmatrix} |_{X_i, Z_i} \sim N.I.D. \left(\begin{pmatrix} \overline{\mathbf{X1}}_i \\ \overline{\mathbf{Z1}}_i \end{pmatrix}, \begin{pmatrix} \sigma_u^2 \iota \iota' + \sigma_\eta^2 I & \sigma_{au} \iota \iota' + \sigma_{\eta\nu} I \\ & \sigma_a^2 \iota \iota' + \sigma_\nu^2 I \end{pmatrix} \right),$$

where $\epsilon_i = \{\epsilon_{i1}, \dots, \epsilon_{iT}\}$ and $e_i = \{e_{i1}, \dots, e_{iJ}\}$. ι is a vector of ones with length T or J , depending on its context. Recall that T represents the total number of observed periods for a given household and J represents the total number of home purchases made by this household.

Appendix 2: Normality of the Distribution

The joint demand model presented here is estimated using maximum likelihood under the assumption of normality. The literature has shown that, relative to conventional GMM estimation, maximum likelihood estimation provides estimates that are less biased, more efficient, and dynamically more stable. This is so because the maximum likelihood estimate exploits the variance-covariance structure of shocks. Indeed, in a

weak identification context, the extent to which the maximum likelihood dominates conventional GMM is striking (Fuhrer, Moore, and Schub (1995)).

These advantages come at the cost of imposing the normality assumption on the random effects. In the current context, the home purchase dummy is used as one of the two dependent variables. Imposing the normal distribution on the purchase dummy is somewhat debatable. When the normality of the errors is relaxed, exclusion restrictions are required to achieve model identification. While the lack of normality does not affect the consistency of the estimates, it does have consequences with respect to model fit, making the validity of some tests doubtful in the presence of high levels of kurtosis. However, the literature has shown there is evidence that even under non-normal distribution assumptions, the tests are asymptotically distributed as chi-square. Based on this, Monte Carlo studies show that maximum likelihood estimates derived from the normality assumption exhibit very little bias even under non-normal conditions (e.g., Bollen (1989)). Therefore, the impact of imposing functional form restrictions on random effects is likely to be small.

Appendix 3: User Cost Construction

The construction of the user cost follows the definition in Poterba (1984). Assume that (i) the cost of borrowing is r^m ; (ii) the opportunity cost of investing equity in the house rather than in some other riskless investment is r^0 ; (iii) houses depreciate at a rate of δ ; (iv) the property tax rate is μ ; (v) the marginal income tax rate is τ ; and (vi) the house value is expected to increase at g . For a given loan-to-value ratio L , the user cost is expressed as a fraction of the real house price, where the fraction is described as:

$$(1 - \tau)[Lr^m + (1 - L)r^0 + \mu] + \delta - g$$

The user cost is stated in after-tax real terms.²² It includes the following components: (a) the after-tax cost of mortgage interest rate, $(1 - \tau)Lr^m$; (b) the after-tax opportunity cost of investing equity in the house rather than in some other riskless assets, $(1 - \tau)(1 - L)r^0$; (c) the after-tax cost of property tax payments, $(1 - \tau)\mu$; (d) the cost of depreciation, δ ; and (e) the expected housing return. g . The risk adjustment component is measured separately in this paper. The expected cost before risk adjustment, UC_{it} , is measured over a five-year period.²³

The components of the user cost are obtained from the following sources:

Expected housing return (g): Imputed from the AR(1)-GARCH(1,1) model in Section 3.2.3.

Opportunity cost of investment (r^0): the U.S. Treasury bill rate.

Mortgage rate (r^m): 30-year conventional mortgage rate from the Federal Home Loan Mortgage Corporation.

Loan-to-value ratio (L): For transactors, this is self-reported. For non-transactors, I assume it to be 70% should a purchase occur.

Property tax rate (μ): The property tax amount comes from a direct survey question in the PSID, except in 1988 and 1989. The wave-VI codebook describes the estimation method for the property tax. Following Flavin and Yamashita (2002), I apply this method and use the information on where a household lives to calculate the relevant property tax.

Marginal income tax rate (τ): I assume the marginal income rate is the sum of the federal tax rate and the state and local tax rates. Following Flavin and Yamashita (2002), the federal tax rate is assumed to be 28%. The state and local tax rates are obtained from income tax data published by Ernst & Young International (the Global Executive, 2000).

Depreciate rate (δ): The depreciation expenses, including the proportional transaction losses associated with future home sales, are assumed to be a fixed 10%.

Appendix 4: Constructing Hedonic-Adjusted Housing Quantity

I construct the hedonic-adjusted housing quantity in two stages. In the first stage, a hedonic-adjusted house price index is constructed for each market in each year. To see how this works, consider the simplest possible case where we compare the prices of houses in two different markets in two different years. The price index of houses in market k in year t is denoted as $HPI_{k,t}$, while the price index of houses in market k' and year s is denoted as $HPI_{k',s}$. The hedonic-adjusted price index is constructed in the following way:

$$\frac{HPI_{k',s}^h}{HPI_{k,t}^h} = \frac{HPI_{k',s}}{HPI_{k',t}} \frac{HPI_{k',t}}{HPI_{k,t}} \quad (\text{A-3})$$

The left-hand side of equation (A-3) represents the hedonic-adjusted house price in market k' in year s normalized by the price of houses in the benchmark market k in the benchmark year t . The right-hand side of equation (A-3) represents a partitioning of this ratio into two parts. The first part, adjusted for any time variations in housing characteristics, compares the house price in year s with the price in year t for market k' . The second part, adjusted for any cross-market variations in housing characteristics, compares house price in market k' with the price in market k in benchmark year t . Empirically, the first part of equation (A-3) is measured by the OFHEO price indices. However, the construction of the second part, the hedonic house price ratio, is much more data-intensive, requiring data not only on house prices in different markets but also on the characteristics of houses in different markets. Fortunately, Malpezzi, Chun, and Green (1998) have computed the hedonic house prices using 1990 Census data. This permits me to compute the second part of equation

(A-3), where the benchmark year is 1990 and the benchmark city is Madison, WI. Following equation (A-3), I then compute quality-adjusted house price indices for each MSA in each given year, where the price index in Madison in 1990 is normalized to be 1.

In the second stage, I obtain a measure of housing quantity by separating housing expenditure into price and quantity components. To see this, define the market value of house j at market k in year t (V_{jkt}) as a product of the physical quantity of this house (H_{jkt}) and the market-level hedonic-adjusted price index (HPI_{kt}^h):

$$V_{jkt} = HPI_{kt}^h \times H_{jkt} \quad (\text{A-4})$$

Taking the log of equation (A-4) yields:

$$\log(H_{jkt}) = \log(V_{jkt}) - \log(HPI_{kt}^h). \quad (\text{A-5})$$

Given equation (A-5), for each transaction reported in the PSID, the physical quantity of the housing stock is imputed using the household's self-reported house value and the hedonic-adjusted price indices derived in the first stage.

Appendix 5: Instrumental Variable Strategy for Price Endogeneity

To control for the potential bias associated with house price variables, I present a test of house price endogeneity using the a control function approach proposed by Blundell and Powell (2004). In order to perform this test, I use lagged changes in construction worker income and lagged changes in the number of construction workers as instruments. These variables are likely to be correlated with house prices through the supply-side relationship but uncorrelated with home purchase decisions. I obtain the construction worker income and employment data from the Current Employment Statistics (CES).

Using the control function approach, I write the predicted housing return and the associated risk as follows:

$$RET_{m,t} = z_{m,t}\pi_1 + \delta_{m,1} + u_{mt}^1, \quad (\text{A-6})$$

$$VOL_{m,t} = z_{m,t}\pi_2 + \delta_{m,2} + u_{mt}^2, \quad (\text{A-7})$$

and $z_{mt} = (z_{mt}^1, z_{mt}^2)$. RET_{mt} and VOL_{mt} are the *predicted* housing return and conditional volatility generated from the AR(1)-GARCH(1,1) model. z_{mt}^1 is a vector of market-specific variables included in the house price regressions, such as average income growth and unemployment rate. z_{mt}^2 is a vector of instruments, including lagged changes in construction worker income and lagged changes in the number of construction workers. δ_m

and u_{mt} are, respectively, market-specific effects and transitory effects.

The control function assumption is that $(\eta_{it}, \nu_{it}) \perp (RET_{mt}, VOL_{mt}) | (\delta_m, z_{mt})$. In order to integrate (η_{it}, ν_{it}) out, I need to know the form of the distribution of (η_{it}, ν_{it}) conditional on u_{mt}^1 and u_{mt}^2 . Assume the joint normality, one could write η_{it} and ν_{it} as linear functions:

$$\eta_{it} = \gamma_1 u_{mt}^1 + \gamma_2 u_{mt}^2 + e_{it}. \quad (\text{A-8})$$

$$\nu_{it} = \gamma_3 u_{mt}^1 + \gamma_4 u_{mt}^2 + v_{it}. \quad (\text{A-9})$$

It follows that the test of endogeneity of housing return and risk variables is a test of the joint significance of the coefficients γ_1 , γ_2 , γ_3 , and γ_4 . In the first stage, I estimate equations (A-6) and (A-7). In the second stage, a linear function of the estimated residuals from the first stage is included in the joint housing demand estimation. That is, I re-estimate equations (1) and (2), where transitory effects are specified as in equations (A-8) and (A-9), with u_{mt}^1 and u_{mt}^2 replaced with their estimated counterparts \hat{u}_{mt}^1 and \hat{u}_{mt}^2 obtained from the first-stage estimation of housing returns and the associated risks.

A chi-square test for joint significance ($\gamma_1 = \gamma_2 = \gamma_3 = \gamma_4 = 0$) takes the value of 1.6063 with a p-value of 0.35. Since it does not appear to be significant, I conclude that the null hypothesis of no endogeneity cannot be rejected.

Notes

¹These changes are based on the Case-Shiller 10-city composite index.

²Flavin and Yamashita (2002) find a low correlation between housing and other financial assets, which suggests that home investment risk cannot be easily diversified. Using S&P/Case-Shiller Home Price Indices, the Chicago Mercantile Exchange (CME) launched housing futures and options for certain MSAs in 2006. However, De Jong, Driessen, and Hemert (2008) find that the economic benefits of having access to these housing futures for homeowners is small. This is mainly because city-level housing futures provide only a limited hedge for idiosyncratic house price risk.

³Cocco (2000) discusses theoretical hedging implications for homeowners, although his work employs numerical simulation techniques rather than empirical estimation methods. Bajari, Chan, Krueger, and Miller (2008) estimate a dynamic structural model that examines the effects of house price decline on homeowners' housing and non-housing consumption. Their work does not consider price risk and hedging implications.

⁴Recent life-cycle models in the household finance literature include Cocco (2005) and Yao and Zhang (2005).

⁵To see why housing stock imbalance affects only the timing decision but not the size decision, consider a household that currently demands a 3,000 square foot house. It is more likely to make a purchase if its existing house is only 1,800 square feet (larger imbalance) than if its existing house is 2,400 square feet (smaller imbalance). In both cases, however, the purchase size decision will be to buy a 3,000 square foot house. Thus, the purchase size decision does not depend on the housing stock imbalance.

⁶Estimates of housing demand models could also suffer from selection bias due to unobserved factors. For example, prudent households are likely to delay home purchase and buy a smaller home in a volatile housing market. Alternatively, a sudden job promotion could prompt households to make an earlier and bigger home purchase regardless of the housing market conditions.

⁷In housing demand models, the initial condition may be treated as exogenous if we assume that a genuinely new home purchase process starts at the time when a household switches from renting to owning. Such an assumption would be more desirable than treating the initial home purchase as coming from a stationary equilibrium, which would be more appropriate if the model focused on frequently purchased consumption goods.

⁸The model is estimated using a sequential quadratic programming (SQP) method written in C. The numerical integration is approximated using the Gaussian-Hermite quadrature method with 16 points. Sixteen points are large enough to allow for some representation of any heterogeneity that might be present in the sample but small enough to be computationally feasible.

⁹Appendix 2 provides a discussion of the benefits and costs of the distributional assumption.

¹⁰The interview started in 1968, was conducted annually until 1997, and took place biennially thereafter. To match the PSID to house price information from the CMHPI, I exclude the period 1968-1979 from the PSID sample.

¹¹For those who stay in a rental house for less than three years between two home purchases, I still consider the two purchases as sequential home purchases and treat the renting period as part of the previous purchase spell.

¹²A household is defined as dissolved if there is a change in the head of the household. In such a case, observations after the change are dropped.

¹³I consider a five-year period as a reasonable forecast horizon – long enough for a household to overcome the transaction cost barriers, yet short enough to be estimated given the time span in the sample. The PSID asks a question as to whether the household would move again in the next couple of years. I do not use this information because the presence of lumpy transaction costs makes it unlikely that households would plan to move again shortly after a home purchase is made.

¹⁴Note that this is different from the measure of hedging incentives in tenure-choice studies. For example, Sinai and Souleles (2005) use the expected length of stay to proxy hedging incentives for

tenure decision makers. For existing homeowners, the longer a household plans to stay in its current residence, the more heavily both financial risk and the housing cost risk will be discounted. Given that both risks are realized at the same time when a future resale occurs, the relative riskiness of the home purchase decision does not depend on the expected length of stay, but rather on the likelihood of moving up and the correlation between the current and future markets.

¹⁵At a more micro level, using transaction data, Goetzmann and Spiegel (1995) estimate housing returns and risks controlling for both temporal and non-temporal variation. This approach is not feasible in the current analysis, since the price data are available only at the market level.

¹⁶Again, I consider five years as a reasonable horizon over which a household makes its forecast about future housing transactions. Note that short- and long-run housing returns and volatility are highly positively correlated.

¹⁷Kleiner, Gay, and Greene (1982) measure the effects on migration patterns of state or local occupational licensing. They consider 14 universally licensed occupations: accountants, architects, engineers, lawyers, dentists, pharmacists, physicians, surveyors, insurance agents, real estate agents, registered nurses, practical nurses, barbers, and cosmetologists. The three-digit occupations are chosen from professional and technical occupations and clerical and service workers. In this paper, the dummy variable ‘licensed occupation’ equals 1 if the head’s occupation belongs to the category specified in Kleiner, Gay, and Greene (1982).

¹⁸This is equivalent to assuming that physical depreciation and appreciation resulting from remodeling or market cycles cancel out on average.

¹⁹An alternative explanation is that households tend to be more reluctant to respond to certain self-predicted changes in their family circumstances. For example, a married couple who anticipate a 30% likelihood of divorce may not want to react to these expectations until the divorce actually occurs.

²⁰In similar spirit, Davidoff (2006) examines the effect of covariance between income risk and house price risk. Sinai and Souleles (2005) evaluate how the share of housing cost in a household’s budget

affects its response to rent volatility.

²¹One might question whether the control group is sufficiently similar to the treatment group where the occupation affects future housing consumption plans. Since the demand regression has controlled for many household characteristics, I argue this assumption holds roughly speaking.

²²While capital gains on housing are taxable in some circumstances, I assume that the bulk of realized capital gains escape taxation due to rollover provisions, the one-time exemption for homeowners over age 55, and the fact that inherited real estate receives a stepped-up cost allowance when transferred. I therefore assume that the capital gains portion is not taxed.

²³The formulation here assumes that all homeowners are itemizers and can take full advantage of the interest and property-tax deductions. Also, it assumes away the variation in marginal tax rates over the 1980-1997 period. The differences in tax provisions for itemizers and non-itemizers and the changes in tax rates over time are worth studying in their own right. However, the advantage of our empirical approach is that the random effect specification would pick up fixed differences in tax treatment across households and that the year dummies would pick up national level changes in tax provisions over time.

References

- Amemiya, T., 1985, *Advanced Econometrics*. Harvard University Press, Princeton NJ.
- Bajari, P., P. Chan, D. Krueger, and D. Miller, 2008, “A Dynamic Structural Model of Housing Demand: Estimation and Policy Implications,” Working Paper, University of Minnesota.
- Banks, J., R. Blundell, Z. Oldfield, and J. Smith, 2004, “House Price Volatility and Housing Ownership Over the Lifecycle,” UCL Discussion Paper.
- Berkovec, J., and D. Fullerton, 1992, “A General Equilibrium Model of Housing, Taxes, and Portfolio Choice,” *Journal of Political Economy*, 100(2), 390–429.
- Berson, D., and L. Henry, 2008, “How Far Will House Prices Decline,” Economic Real Estate Trends, Quarterly Report by PMI Mortgage Insurance Co.
- Blundell, R., and J. Powell, 2004, “Endogeneity in Semiparametric Binary Response Models,” *Review of Economic Studies*, 71(7), 655–679.
- Bollerslev, T., 1986, “Generalized Autoregressive Conditional Heteroscedasticity,” *Journal of Econometrics*, 31(3), 307–327.
- Brueckner, J., 1997, “Consumption and Investment Motives and the Portfolio Choices of Homeowners,” *Journal of Real Estate Finance and Economics*, 15(2), 159–180.
- Caplin, A., S. Chan, C. Freeman, and J. Tracy, 1997, *Housing Partnerships*. MIT Press, Cambridge, MA.
- Case, K., and R. Shiller, 1989, “The Efficiency of the Market for Single-Family Homes,” *American Economic Review*, 79(1), 125–137.
- Cocco, J., 2000, “Hedging House Price Risk with Incomplete Markets,” Mimeo, London Business School.
- , 2005, “Portfolio Choice in the Presence of Housing,” *Review of Financial Studies*, 18(2), 535–567.

- Davidoff, T., 2006, “Labor Income, House Prices, and Homeownership,” *Journal of Urban Economics*, 59(2), 209–235.
- De Jong, F., J. Driessen, and O. V. Hemert, 2008, “Hedging House Price Risk: Portfolio Choice with Housing Futures,” working paper, Stern School of Business, New York University.
- Englund, P., M. Hwang, and J. Quigley, 2002, “Hedging Housing Risk,” *Journal of Real Estate Finance and Economics*, 24(1), 167–200.
- Flavin, M., and S. Nakagawa, 2008, “A Model of Housing in the Presence of Adjustment Costs: A Structural Interpretation of Habit Persistence,” *American Economic Review*, 98(1), 474–495.
- Flavin, M., and T. Yamashita, 2002, “Owner-Occupied Housing and the Composition of the Household Portfolio,” *American Economic Review*, 92(1), 345–362.
- Fuhrer, J., G. Moore, and S. Schub, 1995, “Estimating the Linear-Quadratic Inventory Model: Maximum Likelihood versus Generalized Method of Moments,” *Journal of Monetary Economics*, 35(1), 115–158.
- Goetzmann, W., and M. Spiegel, 1995, “Non-temporal Components of Residential Real Estate Appreciation,” *Review of Economics and Statistics*, 77(1), 199–206.
- , 2002, “The Policy Implications of Portfolio Choice in Underserved Mortgage Markets,” in *Low Income Homeownership*, ed. by N. Retsinas, and E. Belsky. Brookings Institution Press, Washington, D.C.
- Goodwin, T., 1986, “Inflation, Risk, Taxes, and the Demand for Owner-Occupied Housing,” *Review of Economics and Statistics*, 68(2), 197–206.
- Green, R., and K. D. Vandell, 1999, “Giving Households Credit: How Changes in the US Tax Code Could Promote Homeownership,” *Regional Science and Urban Economics*, 29(4), 419–444.
- Grossman, S. J., and G. Laroque, 1990, “Asset Pricing and Optimal Portfolio Choice in the Presence of Illiquid Durable Consumption Goods,” *Econometrica*, 58(1), 25–51.

- Gyourko, J., and P. Linneman, 1996, “An Analysis of the Changing Influences on Traditional Household Ownership Patterns,” *Journal of Urban Economics*, 39(3), 318–341.
- Han, L., 2008, “Hedge House Price Risk in the Presence of Lumpy Transaction Costs,” *Journal of Urban Economics*, 64(2), 270–287.
- Haurin, D., and H. L. Gill, 2002, “The Impact of Transaction Costs and the Expected Length of Stay on Homeownership,” *Journal of Urban Economics*, 51(3), 563–584.
- Heaton, J., and D. Lucas, 2000, “Portfolio Choice and Asset Prices: The Importance of Entrepreneurial Risk,” *Journal of Finance*, 55(3), 1163–1198.
- Henderson, V., and Y. Ioannides, 1983, “A Model of Housing Tenure Choice,” *American Economic Review*, 73(1), 98–113.
- Ioannides, Y., and S. Rosenthal, 1994, “Estimating the Consumption and Investment Demands for Housing and Their Effect on Housing Tenure Status,” *Review of Economics and Statistics*, 76(1), 127–141.
- Kan, K., 2003, “Residential Mobility and Job Changes Under Uncertainty,” *Journal of Urban Economics*, 54(3), 566–586.
- Kleiner, M., R. Gay, and K. Greene, 1982, “Barriers to Labor Migration: The Case of Occupational Licensing,” *Industrial Relations*, 21(3), 383–391.
- Malpezzi, S., G. Chun, and R. K. Green, 1998, “New Place-to-Place Housing Price Indexes for U.S. Metropolitan Areas, and Their Determinants,” *Real Estate Economics*, 26(2), 235–274.
- Mwang, M., and J. Quigley, 2006, “Economic Fundamentals in Local Housing Markets: Evidence from U.S. Metropolitan Regions,” *Journal of Regional Science*, 46(3), 425–453.
- Ortalo-Magne, F., and S. Rady, 2002, “Tenure Choice and the Riskiness of Non-housing Consumption,” *Journal of Housing Economics*, 11(3), 226–279.
- Pashigian, P., 1979, “Occupational Licensing and the Interstate Mobility of Professionals,” *Journal of Law and Economics*, 22(1), 1–25.

- Poterba, J., 1984, "Tax Subsidies to Owner-occupied Housing: An Asset-Market Approach," *Quarterly Journal of Economics*, 99(4), 729–752.
- Rosen, H., K. Rosen, and D. Holtz-Eakin, 1984, "Housing Tenure, Uncertainty, and Taxation," *Review of Economics and Statistics*, 66(3), 405–416.
- Sinai, T., and N. Souleles, 2005, "Owner Occupied Housing as a Hedge Against Rent Risk," *Quarterly Journal of Economics*, 120(2), 763–789.
- , 2008, "Can Owning a Home Hedge the Risk of Moving," Mimeo, the Wharton School, University of Pennsylvania.
- Somerville, T., 1999, "Residential Construction Costs and the Supply of New Housing: Endogeneity and Bias in Construction Cost Indexes," *Journal of Real Estate and Finance Economics*, 18(1), 43–62.
- Spiegel, M., 2001, "Housing Return and Construction Cycles," *Real Estate Economics*, 29(4), 521–551.
- Stein, J., 1995, "Prices and Trading Volume in the Housing Market: A Model with Downpayment Constraints," *Quarterly Journal of Economics*, 110(2), 379–406.
- Tracy, J., and H. Schneider, 2001, "Stocks in the Household Portfolio: A Look Back at the 1990's," *Current Issues in Economics and Finance*, 7(4), 1–6.
- Yao, R., and H. H. Zhang, 2005, "Optimal Consumption and Portfolio Choices with Risky Labor Income and Borrowing Constraints," *Review of Financial Studies*, 18(1), 197–239.

Table 1: Descriptive statistics

Variables	Mean	Sd
house value (in 1982\$)	77,577	264,105
waiting time (in years)	7.53	6.76
age	45.72	18.95
education index	10.92	4.24
family size	2.87	1.44
number of children	1.02	1.18
family income(in 1982\$)	34,515	32,942
white (%)	84.31	36.37
married (%)	79.75	4.02
working status (%)	69.40	46.08
both working (%)	26.91	36.62
licensed occupation (%)	12.15	32.67
self-employed (%)	12.26	32.80
presence of children aged 0-6 (%)	0.10	0.33
presence of children aged 7-12 (%)	0.11	0.35
presence of children aged 13-18(%)	0.32	0.68
number of transactions	3566	
number of observations	37573	
number of households	2567	

The data source is the Panel Study of Income Dynamics (1980-1997). The education index is in the range 1-17 and represents the actual grade of school completed. For example, a value of 8 indicates that the household head completed eighth grade; a value of 17 indicates that the head completed at least some postgraduate work.

Table 2: Housing transaction frequency in the PSID (1980-1997)

	Mean (%)	Sd(%)
Transaction Rate	13.91	33.97
(Trade-Up Moves)/(Total Transactions)	75.53	42.99
(Inter-state Moves)/(Total Transactions)	16.95	37.54
(Inter-MSA Moves)/(Total Transactions)	31.01	54.16

The data source is the Panel Study of Income Dynamics (1980-1997). A move is considered a “trade-up” if the size of the newly purchased house exceeds the size of the previously owned house, where the house size is measured by the deflated gross house value.

Table 3: Summary statistics for conditional housing demand measures

	Mean	Sd
Hedonic-Adjusted House Size	1,395.34	5,357.62
Real House Value (in 1982\$)	77,577.31	26105.5
Square Footage	1,234.71	3,488.04
Number of Rooms	6.15	1.77

The data source is the Panel Study of Income Dynamics (1980-1997). Section 3.2.1 explains how these variables are constructed.

Table 4: Correlation among conditional housing demand measures

Measures	Hedonic-Adjusted Size	Real Value	Square Footage	Number of Rooms
Hedonic-Adjusted Size	1.000			
Real House Value	0.8398*	1.0000		
Square Footage	0.5211*	0.4554*	1.0000	
Number of Rooms	0.5891*	0.5343*	0.8873*	1.0000

This table displays Spearman rank correlation coefficients among alternative measures of conditional housing demand. The correlation coefficients significant at the 1% level or lower are marked with an asterisk.

Table 5: Quarterly housing return and volatility estimates

	mean	sd(mean)
$RET_{m,t-1}$	0.62	0.28
income growth	0.35	0.15
unemployment change	-0.27	0.16
arch1	0.39	0.06
garch1	0.56	0.04

The estimates are cross-markets averages of the AR(1)-GARCH(1,1) model estimated at the MSA level. In addition, sd(mean) indicates standard deviation of the mean estimates. The coefficients on the constants are suppressed.

Table 6: Sequential logit estimation for the hedging propensity

Specification	(I)	(II)	(III)	(IV)
1st level decision: <i>Whether to trade up</i>				
age	-0.37 (0.14)	-0.14 (0.08)	-0.18 (0.12)	-0.09 (0.03)
educ	0.46 (0.12)	0.57 (0.10)	-0.35 (0.12)	-0.34 (0.28)
job tenure	0.36 (0.30)	0.000 (0.19)	0.002 (0.04)	0.03 (0.15)
married	-0.21 (0.12)	-0.27 (0.15)	-0.31 (0.26)	-0.33 (0.16)
family size	-0.25 (0.12)	-0.15 (0.06)	-0.45 (1.28)	-0.57 (0.43)
family income (in log)	-0.11 (0.04)	-0.13 (0.06)	-0.08 (0.11)	-0.12 (0.33)
2nd level decision: <i>Whether to remain in the same MSA, if trading up</i>				
licensed	0.11 (0.05)	0.09 (0.03)	0.07 (0.04)	0.08 (0.05)
“licensed” × job tenure	-0.03 (0.02)	-0.04 (0.03)	0.11 (0.20)	-0.13 (0.08)
presence of children (0-6)	0.02 (0.01)	0.02 (0.01)	0.03 (0.02)	0.06 (0.14)
presence of children (7-12)	0.09 (0.05)	0.07 (0.03)	0.05 (0.02)	0.08 (0.16)
presence of children (13-18)	0.07 (0.03)	0.04 (0.02)	0.03 (0.04)	0.06 (0.03)
dual employment	0.06 (0.02)	0.07 (0.04)	0.06 (0.08)	0.08 (0.05)
self-employed	0.12 (0.04)	0.16 (0.11)	0.18 (0.17)	0.13 (0.08)
age	-0.03 (0.01)	-0.02 (0.01)	-0.02 (0.01)	-0.008 (0.007)
educ	-0.22 (0.06)	-0.18 (0.03)	-0.13 (0.05)	-0.19 (0.03)
family size	0.24 (0.23)	0.12 (0.11)	0.05 (0.10)	-0.02 (0.09)
measures of house size	hedonic adjusted quantity	deflated house value	imputed square footage	number of rooms

Household fixed effects are included in the first level but not in the second level. A move is considered a “trade-up” if the size of the next home purchase exceeds the size of the current home purchase. The size of home purchase is measured by the imputed hedonic-adjusted housing quantity in Specification (I), the deflated house value in Specification (II), the imputed square footage in Specification (III), and the number of rooms in a house in Specification (IV). To conserve space, only selected coefficients are reported. Full results are presented in the web appendix.

Table 7: Mean household characteristics by strength of hedging incentives

Predicted Hedging Propensity	age	educ index	white	male	family income (1982\$)	married	number of children	family size
Low	55	8.26	0.82	0.86	24526	0.82	1.08	3.2
Medium	42	12.28	0.80	0.89	44014	0.85	0.78	3.0
High	37	11.67	0.91	0.78	33355	0.71	0.58	2.5

This is a comparison across households with different levels of predicted hedging propensities, imputed based on the estimates from the first specification in Table 6. Only selected demographic information is reported here. The data source is the PSID (1980-1997). The education index is in the range 1-17 and represents the actual grade of school completed. For example, a value of 8 indicates that the household head completed eighth grade; a value of 17 indicates that the household head completed at least some postgraduate work.

Table 8: Dynamic joint housing demand estimation

Variable	Specification I		Specification II	
	COEF	STD ERR	COEF	STD ERR
<i>Purchase Size Equation</i>				
VOL_{it}	-0.25	(0.08)	-0.21	(0.09)
$VOL_{it} \times HP_{it}$	0.11	(0.05)		
$VOL_{it} \times HP_{it}^m$			0.09	(0.05)
$VOL_{it} \times HP_{it}^h$			0.15	(0.06)
$\ln UC_{it}$	-0.08	(0.04)	-0.08	(0.04)
age	0.002	(0.001)	0.002	(0.001)
age \times age	-0.001	(0.001)	-0.001	(0.001)
educ	0.07	(0.01)	0.07	(0.01)
married	0.22	(0.08)	0.23	(0.08)
family size	0.09	(0.02)	0.08	(0.02)
family income (in logs)	0.32	(0.05)	0.30	(0.05)
<i>Purchase Timing Equation</i>				
VOL_{it}	-0.03	(0.01)	-0.04	(0.02)
$VOL_{it} \times HP_{it}$	0.49	(0.21)		
$VOL_{it} \times HP_{it}^m$			0.52	(0.33)
$VOL_{it} \times HP_{it}^h$			1.64	(0.52)
$\ln UC_{it}$	-0.11	(0.03)	-0.11	(0.03)
waiting time (in years)	-0.13	(0.02)	-0.13	(0.02)
waiting time squared	0.004	(0.001)	0.004	(0.001)
TC_{it} (previous home value)	-0.27	(0.05)	-0.25	(0.05)
Δ income (in logs)	0.04	(0.08)	0.07	(0.08)
Δ married	0.37	(0.16)	0.35	(0.16)
Δ working status	0.28	(0.14)	0.28	(0.15)
increase in # of children	0.12	(0.05)	0.11	(0.05)
decrease in # of children	0.02	(0.04)	0.02	(0.04)
age	-0.006	(0.003)	-0.006	(0.003)
educ	0.09	(0.02)	0.09	(0.02)
<i>Permanent Effects</i>				
σ_a	0.11	(0.05)	0.11	(0.05)
σ_u	0.25	(0.03)	0.24	(0.04)
ρ_{au}	0.28	(0.11)	0.28	(0.11)

Purchase size is measured by the imputed hedonic-adjusted housing quantity. Specification I uses a continuous measure of hedging propensity (HP_{it}). Specification II uses discrete measures of hedging propensity (HP_{it}^m , HP_{it}^h). To conserve space, only selected coefficients are reported. Full results are presented in the web appendix.

Table 9: Marginal effects of housing return and risk

	expected return	price risk (σ)		
		Low (hedging propensity)	Medium	High
<i>Purchase Size</i>	1.75% \uparrow (0.38%)	1.42% \downarrow (0.34%)	0.11% \downarrow (0.05%)	1.06% \uparrow (0.57%)
<i>Purchase Probability</i>	0.93% \uparrow (0.42%)	0.68% \downarrow (0.12%)	0.34% \downarrow (0.21%)	0.45% \uparrow (0.13%)

The numbers are computed based on the coefficients in Specification II from Table 8, where the hedging propensity measures are discrete. Except for the hedging propensity, the effects are evaluated at the sample average. The housing return risk is measured as the square root of the conditional forecast of a five-year housing return. Standard errors are in parentheses.

Table 10: Estimates of the price risk terms from the housing demand model

specification	(I)		(II)		(III)		(IV)		(V)		(VI)	
variable	timing	size	timing	size	timing	size	timing	size	timing	size	timing	size
with continuous measure of HP_{it}												
VOL_{it}	-0.03 (0.01)	-0.25 (0.08)	-0.07 (0.03)	-0.21 (0.09)	-0.02 (0.02)	-0.13 (0.08)	-0.05 (0.04)	-0.18 (0.11)	-0.03 (0.01)	-0.22 (0.10)	-0.03 (0.01)	-0.21 (0.12)
$VOL_{it} \times HP_{it}$	0.49 (0.21)	0.11 (0.05)	0.52 (0.19)	0.13 (0.05)	0.38 (0.27)	0.29 (0.18)	0.10 (0.41)	0.17 (0.06)	0.48 (0.20)	0.10 (0.05)	0.47 (0.24)	0.10 (0.05)
with discrete measures of HP_{it}												
VOL_{it}	-0.04 (0.02)	-0.21 (0.09)	-0.06 (0.01)	-0.18 (0.08)	-0.02 (0.01)	-0.15 (0.14)	-0.01 (0.01)	-0.17 (0.10)	-0.04 (0.02)	-0.22 (0.09)	-0.05 (0.01)	-0.21 (0.12)
$VOL_{it} \times HP_{it}^m$	0.52 (0.33)	0.09 (0.05)	0.63 (0.38)	0.18 (0.12)	0.18 (0.16)	0.02 (0.01)	0.22 (0.18)	0.09 (0.04)	0.46 (0.31)	0.10 (0.06)	0.58 (0.32)	0.13 (0.06)
$VOL_{it} \times HP_{it}^h$	1.64 (0.52)	0.15 (0.06)	1.45 (0.62)	0.21 (0.19)	0.82 (0.51)	0.11 (0.06)	1.33 (0.48)	0.08 (0.05)	1.60 (0.66)	0.13 (0.06)	1.65 (0.63)	0.13 (0.08)
measures of for H_{it} :	hedonic adjusted quantity		deflated house value		imputed square footage		number of rooms		hedonic adjusted quantity		hedonic adjusted quantity	
controls for VOL_{it} :												
interaction with permanent effect	no		no		no		no		yes		yes	
instrument variables	no		no		no		no		no		yes	

The top panel reports estimates from the joint housing demand model where a continuous measure of hedging propensity is used. The bottom panel reports the estimates from the same model where discrete measures of hedging propensity are used. Specification 1 replicates the estimates on the risk terms in Table 8. Specifications 2-4 repeat the estimation in Specification 1 but use different measures of home purchase size. Specification 5 repeats the estimation in Specification 1 but allows a linear dependence between price variables and the household permanent effects. Specification 6 is similar to Specification 1, but with the house price risk variable instrumented by lagged changes in construction worker income and in the number of construction workers. Standard errors are in parentheses.

Table 11: Effects of business ownership on housing demand

variable	timing	size
VOL_{it}	-0.02 (0.01)	-0.22 (0.12)
BO_{it}	0.08 (1.29)	0.17 (0.08)
$VOL_{it} \times HP_{it}$	0.38 (0.15)	0.10 (0.05)
$VOL_{it} \times BO_{it}$	0.87 (0.93)	1.09 (2.45)
$VOL_{it} \times BO_{it} \times HP_{it}$	-0.51 (1.05)	-0.02 (0.91)


Estimates are from a joint housing demand model as specified in Specification I of Table 8, with additional variables on the business ownership dummy and its interaction with price risk and hedging terms. $BO_{it} = 1$ if household i owns a business at time t , and $= 0$ otherwise. HP_{it} is a continuous measure of hedging propensity. For brevity, only selected coefficients are reported. Standard errors are in parentheses.

Table 12: Robustness Checks for Housing Demand Model

specification	(I)		(VII)	
variable	timing	size	timing	size
with continuous measures of HP_{it}				
VOL_{it}	-0.03 (0.01)	-0.25 (0.08)	-0.03 (0.01)	-0.24 (0.08)
$VOL_{it} \times HP_{it}$	0.49 (0.21)	0.11 (0.05)	0.37 (0.16)	0.08 (0.03)
with discrete measures of HP_{it}				
VOL_{it}	-0.04 (0.02)	-0.21 (0.09)	-0.04 (0.02)	-0.21 (0.08)
$VOL_{it} \times HP_{it}^m$	0.52 (0.33)	0.09 (0.05)	0.41 (0.27)	0.08 (0.05)
$VOL_{it} \times HP_{it}^h$	1.64 (0.52)	0.15 (0.06)	1.08 (0.66)	0.11 (0.05)
control age-varying				
risk aversion	no		yes	

The top panel reports estimates from the joint housing demand model where a continuous measure of hedging propensity is used. The second panel reports the estimates from the same model where discrete measures of hedging propensity are used. Specification I replicates the estimates on the risk terms in Table 8. Specification VII repeats the estimation in Specification I, but removes age effects from the predicted hedging propensities. House size is measured by the imputed hedonic-adjusted housing quantity. For brevity, only selected coefficients are reported. Standard errors are in parentheses.

Figure 1: Data structure for an observed household



Household i experiences J_i complete purchase spells. The $(J_i + 1)$ th spell is the last observed spell only for right-censored households. Purchase spell j consists of t_j years. In each year, household i makes two decisions: whether to transact, and the size of the purchase should a transaction occur.

Figure 2: The forecasting tree for constructing hedging incentives

