Hedging House Price Risk With Incomplete Markets

João F. Cocco[1]

First draft: April 1999

This version: September 2000

[1]London Business School, Regent's Park, London NW1 4SA, UK, Tel.: (020) 7262-5050, Fax: (020) 7724-3317. E-mail: jcocco@london.edu. I would like to thank Rui Albuquerque, John Y. Campbell and seminar participants at Harvard, LSE, and the Society for Computational Economics and Finance for comments. Financial support from the Banco de Portugal and Fundação para a Ciência e Tecnologia is gratefully acknowledged.

Abstract

This paper solves a model of the optimal asset and consumption choices of a liquidity constrained investor who derives utility from the consumption of both non-durable consumption goods and housing. Using PSID labor income and house price data I estimate a large positive correlation between income shocks and house price shocks, and a large negative correlation between house prices and interest rates. I use these estimates to parameterize the model. Using the model I evaluate the effects of labor income, interest rate and house price risk on housing choices and investor welfare. Due to the dual role of housing as an asset and a source of consumption services, liquidity constraints are an important determinant of hedging demands.

JEL classification: G12.

Keywords: Hedging demands, asset choices, interest rate risk, house price risk, labor income risk, borrowing constraints.

1 Introduction

This paper studies asset and consumption choices in the presence of illiquid durable consumption goods whose price fluctuates over time (e.g. housing). It is fairly easy to motivate the importance of this topic. Buying a house is the single most important investment that households undertake during their lifetime. Due to the large transaction costs of buying and selling a house there is an important dimension of illiquidity or irreversibility in the home investment. And the price of housing fluctuates considerably over time, and with it the value of the home investment and the wealth of homeowners.

This topic is particularly interesting since the effects of house price risk on consumers' choices are not obvious. The reason is that even though fluctuations in house prices lead to fluctuations in wealth, for an investor with utility defined over housing consumption, homeownership serves as an hedge against fluctuations in the cost of consumption: decreases in the price of housing, and in the wealth of homeowners, tend to be accompanied by a decrease in the implicit rental cost of housing. In a multi-period problem with time-varying house prices investors may seek to hedge their exposure to house price fluctuations by acquiring a different house, and this gives rise to intertemporal hedging demands (Merton, 1971).

The extent to which hedging considerations affect current housing choices depends on the existence of frictions. If markets are complete such that the payoffs of the home investment are spanned by existing financial assets, investors may hedge house price risk using the latter and current consumption choices are not affected. However, in reality it is likely that investors lack the contingent claims that would allow them to perfectly replicate the return-risk characteristics of the home investment. Moral hazard issues and transaction costs make this assumption reasonable (Smith, Rosen and Fallis, 1988), and lead to a dual dimension of the home acquisition decision, involving both investment and consumption considerations.

There are several frictions that are usually of concern to home buyers, including large transaction costs, uninsurable labor income risk and borrowing constraints. Labor income or human capital is undoubtedly an important asset for the majority of households. If markets are complete so that labor income can be capitalized and its risk insured, then human capital can simply be added to current wealth and plays no particular role. But market incompleteness seems to be an important feature to consider. Moral hazard issues prevent investors from borrowing against future labor income, and insurance markets for labor income risk are not well developed. This is the route taken by recent research on portfolio choice (Heaton and Lucas, 1996, 1997).

In this incomplete markets environment I set-up a model of the optimal asset and consumption choices of a finitely lived investor. In each period, an investor endowed with a stream of risky uninsurable labor income has to decide the size of the house to own and how much to consume of non-durable goods. Housing has a dual role, as an asset and a consumption good, and its price fluctuates over time. I allow house price shocks to be correlated with income shocks and interest rates. To capture the illiquid nature of the home investment, I assume a proportional transaction cost of selling a house.

I use data from the Panel Study of Income Dynamics (PSID) to parameterize the labor income and house price processes. This is the largest longitudinal US dataset containing both labor income and housing information. I find that real house prices grew during the sample period, and that their growth rate exhibits positive serial correlation. These features of residential real estate prices are in accordance with the findings of Case and Shiller (1989) and Poterba (1991), among others. Interestingly, I also find that house price shocks are strongly positively correlated with income shocks, and strongly negatively correlated with interest rates.

Using the model I study the effects of labor income risk, interest rate risk, and house price risk on housing investment and investor welfare. I find that both labor income and interest rate risk crowd out housing investment, but due to the highly leveraged nature of investors' portfolios the welfare and portfolio implications of the latter are much larger. I also find that all three sources of risk contribute towards increasing the volatility of the stochastic discount factor. Realistic transaction costs of selling a house are shown to have large effects on asset choices, by reducing the number of house trades and restricting investors' ability to take advantage of serial correlation in house prices. Finally, the characterization of hedging demands for the housing asset emphasizes both horizon effects and the role of liquidity constraints.

There are several related papers, both in the consumption and dynamic asset allocation literatures. The model in this paper builds upon the work of Deaton (1991) and Carroll (1996), who study the effects of labor income risk and borrowing constraints on consumer behavior. Grossman and Laroque (1991) study portfolio choice in the presence of illiquid durable consumption goods, but holding the price of the durable good fixed. A vast literature has emerged studying the effects of return predictability on asset choices (Barberis, 1999, Brennan, Schwarz and Lagnado, 1997, Campbell and Viceira, 1999, Kandel and Stambaugh 1996, Kim and Omberg, 1996, Lynch and Balduzzi, 1999). Return predictability is also a characteristic of the house price process that I consider. Finally, my paper is also related to the literature on optimal investment in the presence of costly reversibility (Bertola and Caballero, 1994, Abel and Eberly, 1996).

This paper is organized as follows. Section 2 presents the model. Section 3 uses PSID data to parameterize the model. Section 4 presents the results. The final section concludes.

2 The Model

2.1 Model Specification

2.1.1 Time Parameters and Preferences

I model the asset and consumption choices of an investor with a time horizon of T periods. The investor derives utility from the consumption of housing and non-durable consumption goods. As Grossman and Laroque (1991) I assume that buying a house is strictly preferred to renting (may be due to tax advantages), so that I do not model the decision to buy versus rent. In each period t, $t = 1, ..., T$, the investor must choose the size of the house to own (H_t), and non-durable goods consumption (C_t). The size of the house should be interpreted broadly as reflecting not only the physical size but also its quality. I assume that the household has isoelastic preferences, generalized to allow for multiple goods. The objective function of household i is:

$$\max_{H_t, C_t} E_0 \sum_{t=1}^{T} \beta^t \frac{\left(H_t^{\delta_t} C_t^{1-\delta_t}\right)^{1-\gamma}}{1-\gamma} + \beta^{T+1} \frac{W_{T+1}^{1-\gamma}}{1-\gamma} \tag{1}$$

where β is the time discount factor, δ_t measures how much the investor values housing consumption relative to other goods consumption, and γ is the coefficient of relative risk aversion. I allow δ_t to vary over time, with household specific characteristics such as family composition (Attanasio and Weber, 1995). The household also derives utility from terminal wealth, W_{T+1}, which can be interpreted as the remaining lifetime utility from reaching age $T+1$ with wealth W_{T+1}. This preference specification allows for precautionary savings (Zeldes, 1989, Deaton, 1991, Carroll, 1996).

2.1.2 Interest Rate Risk

The interest rate describes the state of the economy. I assume that the annual interest rate, r_t, follows a first-order autoregressive $(AR(1))$ process such that:

$$r_t = \mu(1-\phi) + \phi r_{t-1} + \epsilon_t \tag{2}$$

where μ is the mean interest rate, and ϵ_t is a normally distributed white noise shock with mean zero and variance σ_ϵ^2.

2.1.3 Labor Income

The investor is endowed with a stochastic labor income stream, Y_t, $t = 1, ..., T$, against which he cannot borrow. Investor i's age t log labor income, y_{it}, is exogenously given by:

$$y_{it} = f(t, Z_{it}) + u_{it} \tag{3}$$

where $f(t, Z_{it})$ is a deterministic function of age, t, and other individual characteristics, Z_{it}, and u_{it} can be decomposed into an aggregate and idiosyncratic components:

$$u_{it} = \eta_t + \omega_{it}. \tag{4}$$

I assume that idiosyncratic risk is transitory so that ω_{it} is an i.i.d. normally distributed random variable with mean zero and variance σ_ω^2. The aggregate shock, η_t, follows an $AR(1)$ process. To save on state variables I assume that it is perfectly negatively correlated with interest rates such that:

$$\eta_t = -\beta r_t. \tag{5}$$

To assess the validity of this assumption in the parameterization section I use PSID data to estimate the correlation between η_t and r_t.

2.1.4 House Prices and The Housing Investment

There are several sources of risk associated with housing that I capture in the model. Acquiring a house is to some extent an illiquid investment, since the housing market is characterized by large transaction costs. Accordingly, I assume that when selling the house the investor faces a monetary cost equal to a proportion λ of its current market value.

The price of housing fluctuates over time, and with it the value of homeowners' wealth. Let P_t denote the date t price per unit of housing, so that a house of size $\overline{H}$ has price $P_t\overline{H}$ at date t. The price of other goods consumption (the numeraire) is fixed and normalized to one. I assume that the growth rate of house prices, $\Delta p_t = ln(P_t) - ln(P_{t-1})$, follows an $AR(1)$ process.

Frequently, the price of housing in a given area is affected by labor income shocks in the same area. This may be an important source of risk since at times when consumers need cash, the value of the house is also smaller making it difficult to borrow against it. I capture this by assuming that house price growth is perfectly negatively correlated with interest rates (or perfectly positively correlated with aggregate income shocks) such that:[2]

$$\Delta p_t = -\alpha r_t. \tag{6}$$

Thus I assume that when interest rates are low house price growth is high. In the parameterization section below I estimate the correlation between Δp_t and r_t.

It is possible and probably more realistic that house prices follow a more general process,

[2]This assumption greatly simplifies the numerical solution of the problem since it avoids having an additional state variable.

involving a higher-order autoregressive process, or moving average terms (Case and Shiller, 1989, Poterba, 1991). Furthermore, it would also be reasonable to assume that investors do not know the stochastic process governing house prices, the parameters of this process, or even current housing prices. The housing market is a fairly illiquid market, for which price information is not readily available, as it is the case for centralized securities markets. It would be interesting to consider more general processes for house prices, but the introduction of additional state variables would not be computationally tractable.

2.1.5 Borrowing Constraints

The acquisition of real estate is often associated with a leveraged position, which makes the wealth of homeowners' particularly sensitive to house prices. Mortgage contracts, or more generally loan contracts which use the value of real estate as collateral, are often complex. Although it is important to explore the portfolio implications of alternative loan contracts, I simplify the analysis by considering the following type of loan.

Investors may borrow up to the market value of the house they own minus a downpayment. The downpayment is a proportion ψ of the current market value of the house. Let B_t denote borrowing in period t. This leads to the following constraint:

$$B_t \leq (1-\psi)H_tP_t, \ \forall t \ . \tag{7}$$

To avoid having B_t as a choice variable I assume that in each period t constraint (7) holds with equality, as in Carroll and Dunn (1999). Since this assumption imposes that households hold debt that they may be willing to repay, I assume that the date t interest rate on the loan is equal to the date t interest rate on financial savings. In this case the assumption that (7) holds with equality is not restrictive.

While keeping the problem tractable, the type of loan contract considered here has the interesting feature that the borrowing limit is time-varying, and depends both on investor's choices and current house prices. Furthermore, the borrowing limit is, through house prices, correlated with income shocks. Since labor income shocks are positively correlated with house prices the borrowing constraint is more likely to bind at times when borrowing is more valuable.

2.1.6 The Investor's Optimization Problem

The investor starts period t with given levels of housing and financial assets. Then conditional on the state of the economy (the level of interest rates) and current house prices he must choose whether to change houses and the level of other goods consumption. Following Deaton (1991) I denote cash-on-hand in period t by X_t. Let M_t be an indicator variable which takes the value of one when the household moves houses and zero otherwise. The equations describing the evolution of cash-on-hand:

$$X_{t+1} = \tilde{Y}_{t+1} + (X_t - C_t + B_t - B_{t-1})(1 + \tilde{r}_{t+1}) + ((1-\lambda)P_t H_{t-1} - P_t H_t)(1 + \tilde{r}_{t+1})M_t - B_t \tilde{r}_{t+1}, \ \forall t \ . \tag{8}$$

Terminal wealth, W_{T+1} is given by:

$$W_{T+1} = X_{T+1} + (1-\lambda)P_{T+1}H_T - B_T \ . \tag{9}$$

The problem the investor faces is to maximize (1) subject to (3) through (9), plus non-negativity constraints on housing and other consumption. The control variables for the problem are $\{H_t, C_t\}_{t=1}^{T}$. The state variables are $S_t = \{t, X_t, H_{t-1}, r_t, P_t\}_{t=1}^{T}$.

The Bellman equation for this problem is:

$$V(S_t) = \max_{H_t, C_t} \frac{(H_t^{\delta_t} C_t^{1-\delta_t})^{1-\gamma}}{1-\gamma} + E_t \beta^{t+1} V(\tilde{S}_{t+1}) \ , \ t = 1, ..., T. \tag{10}$$

2.1.7 Solution Technique

This problem cannot be solved analytically. The numerical techniques I use for solving it are standard (see Judd, 1998). Given the finite nature of the problem a solution exists and can be obtained by backward induction. I discretize the state space and the choice variables using grids equally spaced in the log scale. The density functions for the random variables were approximated using Gaussian quadrature methods to perform numerical integration (Tauchen and Hussey, 1991). The interest rate process was approximated by a two-state transition probability matrix, so that in each period interest rates can either be low or high. The grid points for this process were chosen using Gaussian quadrature. The choice of a two state transition probability matrix for the interest rate process substantially reduces the dimensionality of the problem since, at any given date t, we only need to keep track of the number of times that interest rates have been high in the past to obtain current house prices.

In period $T+1$ the utility function coincides with the value function. In every period t prior to T, and for each admissible combination of the state variables, I compute the value associated with each combination of the choice variables. This value is equal to current utility plus the expected discounted continuation value. To compute this continuation value for points which do not lie on the grid I use cubic spline interpolation. The combinations of the choice variables ruled out by the constraints of the problem are given a very large (negative) utility such that they are never optimal. I optimize over the different choices using grid search.

3 Parameterization

The objective of this section is to estimate the parameters of the interest rate, house prices and labor income stochastic processes. The house price and income data used is from the family questionnaire of the Panel Study of Income Dynamics (PSID) for the years 1970 through 1992. This is the largest longitudinal US dataset containing both labor income and housing information. Families that were part of the Survey of Economic Opportunities subsample were dropped to obtain a random sample.

3.1 Interest Rate

To estimate equation (2) I use as a measure of the interest rate the 6-month real T-bill rate from 1970 to 1992 (the time period for which I also have labor income and housing information). Table 1 shows the results.

3.2 Labor Income

I use a broad definition of labor income so as to implicitly allow for insurance mechanisms - other than asset accumulation - that households use to protect themselves against pure labor income risk. Labor income is defined as total reported labor income plus unemployment compensation, workers compensation, social security, supplemental social security, other welfare, child support and total transfers (mainly help from relatives), all this for both head of household and if present his spouse. Observations which still reported zero for this broad income category were dropped. Labor income defined this way is deflated using the Consumer Price Index, with 1992 as the base year. I drop observations in the upper and bottom 2% percentiles of the distribution of real income changes.

The estimation controls for family-specific fixed effects. The function $f(t, Z_{it})$ is assumed to be additively separable in t and Z_{it}. The vector Z_{it} of personal characteristics, other than age and the fixed household effect, includes marital status, and household composition. Household composition equals the additional number of family members in the household besides the head and (if present) spouse. Figure 1 shows the fit of a second order polynomial to the estimated age dummies. This is the age profile passed on to the simulation exercise. I focus on the behavior of young liquidity constrained households so that I truncate the problem at age forty-five.

The residuals obtained from the fixed-effects regressions of (log) labor income on $f(t, Z_{it})$ can be used to estimate σ_η^2 and σ_ω^2. Define Y_t^* as:

$$\log(Y_{it}^*) \equiv \log(Y_{it}) - \hat{f}(t, Z_{it}). \tag{11}$$

Using (4) to substitute out gives:

$$\log(Y_{it}^*) = \hat{\eta}_t + \hat{\omega}_{it} \tag{12}$$

Averaging across individuals gives:

$$\overline{\log(Y_{it}^*)} = \hat{\eta}_t. \tag{13}$$

The variance of $\hat{\eta}_t$ is obtained immediately as the variance of $\overline{\log(Y_t^*)}$. Subtracting this variance from the variance of $\hat{\varepsilon}_{it}$ gives the variance of $\hat{\omega}_{it}$. The estimated coefficients are presented in table 1. The standard deviation of temporary labor income shocks is as large as 0.38, although there probably is substantial measurement error that biases our estimate upwards. Therefore in the simulation exercise I have decided to use a smaller value, or σ_ω equal to 0.20.

The standard deviation of aggregate labor income shocks is much smaller, and equal to 0.034. Aggregate labor income shocks are strongly negatively correlated with interest rates:

the coefficient of correlation is as large as -0.45 and statistically significant at the 5% level (Table 2). This value is below minus one, but is large and negative.

3.3 House Prices

Homeowners in the PSID are asked to assess the current (at the date of the interview) market value of their house.[3] Therefore the market value of the house does not correspond to a real transaction. A major concern with self-assessed values is that households, when asked about the current market value of their house, do not try to rationally assess this value. However, Skinner (1994) compared the self assessed house values in the PSID to the objective measures of the Commerce Department, and found that the two series are quite close in mapping housing price changes in the 1970's and 1980's.

The self assessed value of the house was deflated using the Consumer Price Index, with 1992 as the base year, to obtain real house prices. I drop all observations corresponding to families who moved since the previous interview. The value of the house in the PSID is truncated at 99999 US dollars prior to 1975, and at 999999 US dollars after this year. I drop truncated observations. I also drop observations in the upper and bottom 2 percentiles of the distribution of real house price changes.

In the model the price of different houses are perfectly correlated. This clearly is a limitation of the analysis, but it is also a limiting case of the desirability of housing as an hedge. Define $p_{it} \equiv \log(P_{it})$ where P_{it} is the real price of house i at time t. Averaging across houses I obtain for each year t an index of house prices:[4]

[3] The current market value of the house refers to the month of March whereas labor income is from January to December.

[4] By averaging across the price of different houses I am mainly concerned with common house price risk. For example, if we assume that the price of house i at time t, p_{it}, is given by $p_{it} = p_t + {}^{2}{}_{it}$ where ${}^{2}{}_{it}$ is a purely idiosyncratic i.i.d. shock with zero mean, by averaging across i we obtain an estimate of the common

$$p_t = \frac{\sum_{i=1}^{N_t} p_{it}}{N_t}, \quad t = 1970, ..., 1992. \tag{14}$$

where N_t is the number of observations at time t. Δp_t is obtained by taking first differences. Figure 2 plots the evolution over time of p_t. Real house prices grew during the sample period particularly in the 1970's. Mankiw and Weil (1989) linked this increase in house prices to the changing demographic structure of the US population, although there has been considerable debate in the literature as to whether this was the main reason of the increase, and whether it could have been anticipated. The average house price increase over the sample period is as high as 1.4 percent (Table 1). The growth rate of house prices is fairly volatile, with a standard deviation of 0.045, although this volatility is lower than if idiosyncratic house price risk is also considered.

Table 2 shows the correlation between Δp_t, aggregate income shocks, and interest rates. House price growth is strongly positively correlated with aggregate income shocks: the coefficient of correlation is equal to 0.79 and significant at the one percent level. House price growth is also negatively correlated with interest rates as assumed in the model. Although the coefficient of correlation is significantly below minus one, it is as large as -0.39.

3.4 Other Parameters

I focus on the behavior of young liquidity constrained households, so that I truncate the problem at age 45, and set T equal to 20. In the baseline case I assume a time discount factor equal to 0.98, a coefficient of relative risk aversion equal to 3, and δ_t constant and equal to 0.10. This value implies that investors choose a house which on average is roughly three times their current labor income.

component in house prices. One important source of house price risk not captured by the model is idyosincratic risk.

The downpayment constraint is equal to 15 percent of the current market value of the house. Smith, Rosen, and Fallis (1998) estimate the monetary component of the transaction costs of changing house to be approximately 8-10 percent of the unit being exchanged. This estimate comprises transaction costs associated with search, legal costs, costs of readjusting home furnishings to a new house, and a psychic cost from disruption. Accordingly I set λ equal to 8%. Table 3 summarizes these parameters.

4 Results

4.1 Benchmark Results

To study the behavior of the variables in the model, I calculate cross-sectional averages across two thousand households receiving different draws of income and asset returns and plot them against age. Figure 4 plots labor income, housing, non-durable consumption, and financial savings. The average consumer is borrowing constrained. Non-durable consumption tracks income very closely early in life. The little financial savings that accumulate are due to a precautionary savings motive.

The cross-sectional averages plotted in Figure 4 hide considerable variability at the household level. Table 4 shows several summary statistics which allow us to study the nature of this variability. From the work of Deaton (1991) and Carroll (1996) we know that a prudent consumer, facing borrowing constraints and income risk, accumulates a buffer-stock of assets that allows him to smooth consumption over time, by saving in good times and dissaving in bad times, or when faced with negative income shocks. The first and second rows of Table 4 show the average and standard deviation of non-durable consumption growth, that is, of $\Delta c_t = ln(C_t) - ln(C_{t-1})$. The table shows both unconditional and conditional values (conditional on the state of the economy, or the level of interest rates). Unconditionally, non-durable consumption grows at an average rate of 4.52 percent. But this growth rate is as high as ten percent when interest rates are low (and house prices and labor income are high), and as low as minus one percent when interest rates are high.

Households are able to some extent use financial assets as a buffer-stock. The unconditional standard deviation of consumption growth is equal to 0.142, which is significantly lower than the standard deviation of income growth. As for the growth rate, this unconditional value does not let us see the whole picture: the standard deviation of consumption growth is higher when interest rates are high, so that households are less successful at using assets as a buffer-stock in bad states of the world.

Households' ability to use assets as a buffer-stock is important since consumption volatility plays an important role in many asset pricing models. As it is well known the consumption capital asset pricing model is often rejected because aggregate non-durable consumption is too smooth. Indeed, Hansen and Jagannathan (1991) have shown that a volatile stochastic discount factor is needed to explain a high equity premium. Although my model is set in partial equilibrium and therefore cannot be used to price assets, it may be used to study how far it goes towards generating a volatile stochastic discount factor. The last row of Table 4 shows the standard deviation of the ratio of marginal utility of non-durable consumption in consecutive periods. The model is able to generate a volatile stochastic discount factor, with an unconditional standard deviation of 0.394. In addition there is substantial time-variation in the volatility of the stochastic discount, with a standard deviation equal to 0.206 when interest rates are low, but as high as 0.461 in bad states of the world. Remember that bad states of the world differ from good states along three dimensions: high interest rates, low permanent income, and low house price growth. In the next subsection I evaluate the contribution of income risk, interest rate risk, and house price risk towards the time-variation in the volatility of the stochastic discount factor.

Table 4 also characterizes house trades for the benchmark parameters. The large transaction costs of selling a house limit the number of house trades. Overall there are over 2100 house trades, so that each household on average trades houses 1.09 times or once every 9.58 years. Most households who trade houses do so to increase the size of the house they own, as result of borrowing constraints and an increasing labor income profile. In addition most households (roughly three quarters) trade houses when interest rates are low. When interest rates are low and, due to positive serial correlation, are expected to remain low, expected interest payments are low, labor income is high, and house price growth is high. Therefore, when interest rates are low households have more resources available and housing is a more attractive investment.

Although most households trade houses to increase size, roughly five percent of them (116 households) do so to decrease house size. These are households who are faced with high interest rates and low permanent income and cannot afford their current house. These households contribute significantly towards the higher volatility of the stochastic discount factor in the high interest rates state of the world.

4.2 Evaluating The Effects of Income Risk, Interest Rate Risk, and House Price Risk

In my model there are three types of shocks: income, interest rate, and house price shocks. To evaluate their effects on asset choices I solve the model setting the variance of aggregate income shocks, interest rate shocks, and house price shocks to zero.

The first panel of Table 5 shows the effects of setting the variance of aggregate income shocks to zero ($\sigma_\eta = 0$). Labor income risk crowds out housing investment: on average, in the absence of aggregate income risk, households buy houses that are 2.2 percent larger. As in models of portfolio choice with risky financial assets and nontradable labor income (e.g. Heaton and Lucas, 1997), labor income risk is a source of background risk. It also leads to higher financial asset holdings. The larger housing investment and the lower financial asset holdings explain why, in the absence of aggregate income shocks, households do not do a much better job at smoothing non-durable consumption than in the benchmark case.

The last row shows the welfare effects of setting the variance of aggregate income shocks to zero. To compute such a measure I evaluate the utility associated with this scenario by computing the expectation of discounted lifetime utility at the initial date. I then renormalize discounted utility into consumption-equivalent units, and compute the percentage difference between the no aggregate income risk scenario and the benchmark case. As expected, households are better off in the absence of aggregate income shocks: consumption would have to be 0.58 percent higher at every date in the benchmark case to ensure the same level of lifetime expected utility.

The effects of eliminating interest rate risk are much larger than those of eliminating aggregate income risk, even though the variance of the latter is larger (Table 5, $\sigma_\epsilon = 0$ scenario). In the no interest rate risk scenario the interest rate is constant and equal to its unconditional mean, so that low and high interest rate states now differ only on the level of aggregate income and the growth rate of house prices. Interest rate risk crowds out housing investment: in the absence of interest rate risk average house size is 13.2 percent larger than in the benchmark case. In addition average financial asset holdings are 23.5 percent lower. Eliminating interest risk also has large welfare effects: consumption would have to be 1.63 percent higher in the benchmark case to ensure the same level of lifetime expected utility. Since borrowing-constrained households hold a highly leveraged portfolio, interest rate risk is a major source of background risk.

The final panel of table 5 shows the effects of setting the variance of house price shocks to zero. In this last scenario house prices are fixed and equal to one. Investors are worse off than in the benchmark case. Of course this result is not due to the elimination of house price risk, but to the fact that in the benchmark case housing is an asset with positive average return. When house prices are fixed average house size is larger. The main reason for this is that when households trade houses, they tend to increase house size by more: the average house size increase for those who trade houses is 35.43, compared to only 30.21 in the benchmark case. In the benchmark case, due to the positive expected return on housing, house prices tend grow over time. This means that housing consumption becomes more expensive relative to non-durable consumption, and there is a shift in consumption towards non-durable goods. This result is interesting since it highlights the dual role of housing in my model, as both an asset and a source of consumption services.

4.3 Evaluating The Effects of Transaction Costs

The effects of transaction costs of adjusting stockholdings have been studied by Constantinides (1986), Davis and Norman (1990), Heaton and Lucas (1997), Lynch and Balduzzi (1999), among others. The results can be loosely summarized as follows. When investors trade for the sole purpose of portfolio re-balancing, transaction costs have a small effect on portfolio allocation as in the presence of transaction costs investors choose to trade infrequently. When investors trade for the purpose of consumption smoothing as well, transaction costs have large effects as investors choose to trade frequently, even though there are costs of doing so. Transaction costs of adjusting the level of housing have potentially large effects on asset and consumption choices. Investors trade houses for the purpose of consumption smoothing as well as portfolio rebalancing. And transaction costs of adjusting the level of housing are substantial. A value equal to ten percent of the unit being exchanged is considered reasonable by Smith, Rosen, and Fallis (1988).

Table 6 shows the effects of setting λ equal to 0.6, 0.8 (the benchmark value), and 0.10. Transaction costs substantially reduce the number of house trades. Whereas in the benchmark case investors on average trade houses once every 9.58 years, when transaction costs are lower they do so once every 5.20 years. This higher number of trades is mainly the result of investors trying to take advantage of serial correlation in interest rates and house price growth, by acquiring bigger houses when interest rates are low, and trading these for smaller houses when interest rates increase. The result is an average house size growth of 17.5 percent when interest rates are low, and -8.73 percent when interest rates are high. The corresponding values for the benchmark case are 8.29 percent and 1.81 percent, respectively, so that investors behave less aggressively when transaction costs are higher.

Unconditionally, transaction costs of changing the level of housing crowd out housing investment, as do income risk and interest rate risk. However, while the elimination of interest rate (or income) risk leads to a higher average level of housing in both low and high interest rate states, lowering transaction costs increases housing in low interest rate states, but may decrease it in high interest rate states (as in the λ equal to 6 percent scenario compared with the benchmark case). Of course this is the result of a more aggressive market timing in the low transaction costs scenario. As expected households are better off in welfare terms when transaction costs are lower.

4.4 Risk Aversion

Table 7 shows the results for different values of the coefficient of relative risk aversion. As in models of portfolio choice in the presence of stock return predictability (e.g. Campbell and Viceira, 1999, Lynch and Balduzzi, 1999) I also find that less risk averse investors are more aggressive in their market timing. This shows up in the difference between average log house price growth in low and high interest rate states. For γ equal to 5 average log house price growth is 7.1 percent in low interest rate or high expected house price growth states, and 2.4 percent in high interest rate or low expected house price growth states. Instead, for γ equal to 3 average log house price growth is higher and equal to 8.3 percent when interest rates are low, and lower and equal to 1.8 percent when interest rates are low. That is, low risk γ investors increase house size by more when expected returns are positive, and by less when interest rates are high. Thus low γ investors tend to trade houses and incur the associated transaction costs more often. Unconditionally, low γ investors also hold on average a longer position on the housing asset: the percentage house size difference between $\gamma = 3$ and $\gamma = 5$ investors is as large as 27 percent.

Whereas all these qualitative results might be expected from the recent literature on portfolio choice in the presence of stock return predictability, the large quantitative impact of γ is somewhat surprising. This is because for power utility γ plays a dual role, measuring both risk aversion and elasticity of intertemporal substitution. Less risk averse investors are willing to take more risk. In addition, low γ investors are more willing to substitute consumption intertemporally in response to the state of the economy. That is, low γ investors are more willing to adjust their consumption decisions to take advantage of low interest rates or a high expected growth rate of house prices. This is important in a model with housing since to follow a more aggressive portfolio rule, i.e. to buy a larger house when house prices are expected to grow investors need to save a larger down-payment. The dual role of γ, measuring both risk aversion and elasticity of intertemporal substitution, is also important for understanding the hedging demands that the model generates, which I characterize in the next section.

4.5 Hedging Demands

In my model housing is both an asset and a source of consumption services. When at the initial date an investor acquires a house of a given size, he also acquires a hedge against fluctuations in the price of housing. But since most investors later on trade houses to increase house size, the house they acquire at the initial date is only a partial hedge against fluctuations in the price of housing. In this section I investigate how current asset and consumption choices are affected by the desire to hedge higher expected future housing consumption needs.

In order to buy a larger house to hedge higher future housing consumption needs investors have to save a larger down-payment. For this reason the extent to which investors are borrowing constrained is an important determinant of current housing choices and hedging demands.

In particular, borrowing constraints explain why in the benchmark case investors do not buy a larger house early on, even though they expect to increase house size in the future. Similarly, consider two borrowing-constrained investors with the same initial income and financial assets, but facing different expected growth rates of labor income. A house of a given size is less adequate as an hedge for the investor who faces a steeper labor income profile since he expects to increase house size by more in the future. But does this mean that this investor will buy a bigger house at the initial date? Not necessarily. The steeper labor income profile makes borrowing constraints more severe, and reduces investors' willingness to save the larger down-payment required. Thus investors face a trade-off between buying a bigger house which is a better hedge and buying a smaller house to relax liquidity constraints.

Perhaps the cleanliest way to investigate the nature and the determinants of this trade-off is through preference parameter shifts, namely shifts in δ_t. Recall that δ_t measures preference for housing relative to non-durable consumption. In the benchmark case it is equal to 0.10 for all t. I first consider the case of δ_t equal to 0.10 for $t < 34$, but equal to 0.15 after this date. Therefore, for all periods $t < 34$ instantaneous utility is the same as in the benchmark case, but the investor expects a future preference tilt towards housing consumption. Figure 5 plots cross-sectional averages of house size across two thousand households, and compares the results to the benchmark case.

Surprisingly, at the initial date investors on average buy smaller houses than in the benchmark case. The reason is that investors expect to trade houses earlier, and therefore at the initial date choose to buy a smaller house in order to relax liquidity constraints. Thus, even though due to the preference tilt towards housing consumption $\delta_t = 0.15$ for $t \geq 34$ investors are short housing by more than $\delta_t = 0.10\ \forall t$ investors, they choose to buy an asset which is less adequate as an hedge. It is only when the horizon shortens, and investors trade houses, that average house size becomes larger than in the benchmark case. The importance of liquidity constraints can be assessed by considering housing choices when the constraints are relaxed. One way of doing so is by endowing investors with positive initial wealth, as when initial wealth is sufficiently large borrowing constraints are no longer binding. Table 8 shows the results for initial wealth, W_0, equal to ten thousand dollars. In this case, and at all horizons, investors who are shorter housing buy on average larger houses.

Table 8 also shows the results for γ equal to five. For higher γ hedging demands are larger: the percentage difference in average house size between $\delta_t = 0.15$ for $t \geq 34$ and $\delta_t = 0.10$ $\forall t$ is at all horizons larger than when γ equal to three. Furthermore, due to the larger hedging demands, $\gamma = 5$ investors are not willing to at the initial date buy smaller houses to relax liquidity constraints.

5 Conclusion

This paper characterized the optimal asset and consumption choices of a liquidity-constrained investor who derives utility from the consumption of both housing and non-durable consumption goods. Using PSID labor income and house price data I estimated a large positive correlation between income shocks and house price shocks, and a large negative correlation between house prices and interest rates. These estimates were used to parameterize the model such that when interest rates are low, labor income is high, and house prices are growing fast.

Using the model I studied the effects of labor income risk, interest rate risk, and house price risk on housing investment. I found that both labor income and interest rate risk crowd out housing investment, but due to the highly leveraged nature of investors' portfolios the welfare and portfolio implications of the latter are much larger. I also found that all three sources of risk contribute towards increasing the volatility of the stochastic discount factor. Realistic transaction costs of selling a house were shown to have large effects on asset choices, by reducing the number of house trades and restricting investors' ability to take advantage of serial correlation in house prices. The characterization of hedging demands for the housing asset emphasized the role of liquidity constraints. Liquidity constraints have such an important role due to the dual dimension of housing as both an asset and a source of consumption services.

There are several limitations of the analysis which are important to pursue in future research. First, I have considered a very particular type of debt contract. Given the whole array of debt contracts that investors have available, and the financial importance of the mortgage decision, it is important to study the portfolio and welfare implications of other debt contracts. Second, in my model investors have access to a single financial asset, which can be interpreted as short-term treasury bills. In reality investors have access to other financial assets, most importantly long-term bonds and equities. An important step in future research is to study how portfolio allocations to bonds and equities interact with house price risk and housing choices. Finally, I have assumed that the parameters of the house price process were known. Given the large standard deviations associated with these estimates, it is important to study the effects of parameter uncertainty on hedging demands, perhaps in a Bayesian setting.

References

[1] Abel, A. and J. Eberly (1996), "Optimal Investment With Costly Reversibility," *Review of Economic Studies*, 63, 581-593.

[2] Barberis, N. (1999), "Investing for The Long Run When Returns are Predictable," *Journal of Finance*.

[3] Bertola, G. and R. Caballero (1994), "Irreversibility and Aggregate Investment," *Review of Economic Studies*, 61, 223-246.

[4] Brennan, M. , E. Schwartz and R. Lagnado (1997), "Strategic Asset Allocation," *Journal of Economic Dynamics and Control*, 21, 1377-1403.

[5] Campbell, J. Y. and L. Viceira, (1999), "Consumption and Portfolio Decisions When Expected Returns are Time-Varying," *Quarterly Journal of Economics*, 114, 433-495.

[6] Carroll, C. D. (1997), "Buffer-Stock Saving and the Life-Cycle/Permanent Income Hypothesis," *Quarterly Journal of Economics*, 112, 1-55.

[7] Carroll, C. D. and W. Dunn (1997), "Unemployment Expectations, Jumping (S,s) Triggers, and Household Balance Sheets," *Macroeconomics Annual*, ed. by B. Bernanke and J. Rotemberg, 165-229, MIT Press, Cambridge MA.

[8] Case, K. E., and R. J. Shiller (1989), "The Efficiency of the Market for Single-Family Homes," *American Economic Review*, 79, 125-137.

[9] Constantinides, G. M., (1986), "Capital Market Equilibrium With Transaction Costs," *Journal of Political Economy*, 94, 864-862.

[10] Davis, M., and A. Norman (1990), "Portfolio Selection With Transaction Costs," *Mathematics of Operations Research* 15, 676-713.

[11] Deaton, A. (1991), "Savings and Liquidity Constraints," Econometrica, 59, 1221-1248.

[12] Grossman, S. J. and G. Laroque (1991), "Asset Pricing and Optimal Portfolio Choice in The Presence of Illiquid Durable Consumption Goods," Econometrica, 58, 25-51.

[13] Heaton, J. and D. Lucas, (1996), "Evaluating The Effects of Incomplere Markets on Risk Sharing and Asset Prices," Journal of Political Economy 104, 443-487.

[14] Heaton, J. and D. J. Lucas (1997), "Market Frictions, Saving Behavior and Portfolio Choice," Macroeconomic Dynamics, 1, 76-101.

[15] Judd, K. L. (1998), Numerical Methods in Economics, The MIT Press, Cambridge, Massachusetts.

[16] Kandel, S. and R. Stambaugh (1996), "On the Predictability of Stock Returns," Journal of Finance, 51, 385-424.

[17] Kim, T. S. and E. Omberg (1996), "Dynamic Nonmyopic Portfolio Behavior," Review of Financial Studies, 9, 141-161.

[18] Lynch, A. W. and P. Balduzzi (1999), "Predictability and Transaction Costs: The Impact on Rebalancing Rules and Behavior," Journal of Finance forthcoming.

[19] Mankiw, N. G. and D. Weil (1989) "The Baby Boom, the Baby Bust, and the Housing Market," Regional Science and Urban Economics 19, 1989, 235-258.

[20] Merton, R. (1971), "Optimum Consumption and Portfolio Rules in a Continuous-Time Model," Journal of Economic Theory 3, 373-413.

[21] Poterba, J. M. (1991), "House Price Dynamics: The Role of Tax Policy and Demography," Brookings Papers on Economic Activity, 2, 143-83.

[22] Skinner, J. (1994), "Housing and Saving in the United States," in Y. Noguchi and J. M. Poterba, eds., Housing Markets in the United States and Japan, University of Chicago Press, Chicago, 191-213.

[23] Smith, L. B., K. T. Rosen, and G. Fallis (1988), "Recent Developments in Economic Models of Housing Markets," Journal of Economic Literature, Vol. XXVI, 29-64.

[24] Tauchen, G. and R. Hussey (1991), "Quadrature-Based Methods for Obtaining Approximate Solutions to Nonlinear Asset Pricing Models," Econometrica, 59, 371-396.

[25] Zeldes, S., (1989), "Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence," Quarterly Journal of Economics, 104, 275-298.

Table 1: Estimated parameters of the interest rate, house price, and labor income processes. The interest rate measure is the 6-month real T-bill rate from 1970 to 1992. μ is the mean interest rate, ϕ is the estimated first-order autoregressive coefficient, and σ_ϵ is the standard deviation of the residuals. σ_ω and σ_η are the standard deviations of idyosincratic and aggregate labor income shocks estimated using household-level labor income data from the PSID. $\overline{\Delta p_t} = \overline{ln(P_t) - ln(P_t)}$ is average house price growth, where P_t is an index of house prices in year t constructed using PSID data. $\sigma(\Delta p_t)$ is the standard deviation of the growth rate of house prices. The labor income and house price data is from 1970 to 1992.

Parameter	Value
μ	0.014
ϕ	0.775
σ_ϵ	0.016
σ_ω	0.381
σ_η	0.034
$\overline{\Delta p_t}$	0.014
$\sigma(\Delta p_t)$	0.045

Table 2: Correlation Matrix and p-values. This table shows the correlation coefficients between the growth rate of house prices, Δp_t, aggregate labor income shocks, η_t, and real interest rates. The labor income and house price data is from the PSID. The interest rate measure is the 6-month real T-bill rate. All data is annual from 1970 to 1992. P-vlaues are shown below.

	Δp_t	η_t	r_t
Δp_t	1.000		
η_t	0.786	1.000	
	0.000		
r_t	-0.390	-0.449	1.000
	0.080	0.032	

Table 3: Other Parameters. This table shows the model parameters used in the benchmark case.

Description	Parameter	Value
Time Horizon	T	20
Discount factor	β	0.98
Risk aversion	γ	3
Preference for housing	δ	0.10
Down payment	ψ	0.15
Transaction cost	λ	0.08

Table 4: Benchmark Results. This table shows the results obtained from simulating the model with the parameters shown in Tables 1 and 3. $\overline{\Delta c_t} = \overline{ln(C_t) - ln(C_{t-1})}$ is average non-durable consumption growth and $\overline{\Delta h_t}$ is average house consumption growth. $\sigma(\Delta c_t)$ and $\sigma(\Delta h_t)$ are the standard deviations of non-durable consumption and house size growth, respectively. $\overline{H_t}$ is average house size. $\overline{P_t H_t}$ is average house value. $\overline{Fin.Assets}$ is average financial savings. The table also shows the total number of house trades, as well as the number of house trades associated with an increase ($H_t > H_{t-1}$) and a decrease ($H_t < H_{t-1}$) in house size. $\sigma(SDF)$ is the volatility of the stochastic discount factor. The table shows results conditional on the level of interest rates and unconditional values.

	r_t low	r_t high	Uncond.
$\overline{\Delta c_t}$	0.100	-0.010	0.045
$\sigma(\Delta c_t)$	0.115	0.147	0.142
$\overline{\Delta h_t}$	0.083	0.018	0.051
$\sigma(\Delta h_t)$	0.285	0.201	0.249
$\overline{H_t}$	54.300	44.719	49.528
$\overline{P_t H_t}$	71.247	48.189	59.763
$\overline{Fin.Assets}$	10.122	4.438	7.291
# House Trades	1576	598	2174
$\#H_t > H_{t-1}$	1576	482	2058
$\#H_t < H_{t-1}$	0	116	116
$\sigma(SDF)$	0.206	0.461	0.394

Table 5: The Effects of Income, Interest Rate, and House Price Risk. This table shows the results obtained from simulating the model setting the variance of aggregate income shocks, interest rate shocks, and house price shocks to zero. $\overline{\Delta c_t}$ and $\overline{\Delta h_t}$ are average non-durable consumption and house consumption growth. $\sigma(\Delta c_t)$ and $\sigma(\Delta h_t)$ are the respective standard deviations. $\overline{H_t}$ is average house size. $\overline{P_t H_t}$ is average house value. $\overline{Fin.Assets}$ is average financial savings. The table also shows the total number of house trades, as well as the number of house trades associated with an increase $(H_t > H_{t-1})$ and a decrease $(H_t < H_{t-1})$ in house size. $\sigma(SDF)$ is the volatility of the stochastic discount factor. Welfare measures the percentage increase in consumption in the benchmark case needed to ensure the same level of lifetime utility. The table shows results conditional on the level of interest rates and unconditional values.

	$\sigma_\eta = 0$			$\sigma_\epsilon = 0$			$\sigma(\Delta p_t) = 0$		
	r_t low	r_t high	Uncond.	r_t low	r_t high	Uncond.	r_t low	r_t high	Uncond.
$\overline{\Delta c_t}$	0.093	-0.001	0.046	0.090	-0.000	0.045	0.062	0.014	0.038
$\sigma(\Delta c_t)$	0.113	0.151	0.141	0.108	0.143	0.135	0.102	0.128	0.118
$\overline{\Delta h_t}$	0.090	0.006	0.048	0.081	0.032	0.056	0.084	0.032	0.058
$\sigma(\Delta h_t)$	0.291	0.1936	0.251	0.283	0.207	0.249	0.291	0.195	0.250
$\overline{H_t}$	56.014	45.211	50.633	60.213	51.884	56.064	56.613	49.065	52.854
$\overline{P_t H_t}$	73.587	48.178	60.932	78.774	55.526	67.195	56.613	49.065	52.854
$\overline{Fin.Assets}$	8.458	4.172	6.323	7.576	3.560	5.576	4.106	3.207	3.658
# House Trades	1781	627	2408	1576	625	2201	1668	501	2169
$\#H_t > H_{t-1}$	1781	426	2207	1576	594	2170	1668	501	2169
$\#H_t < H_{t-1}$	0	201	201	0	31	31	0	0	0
$\sigma(SDF)$	0.207	0.457	0.385	0.203	0.428	0.363	0.215	0.352	0.302
Welfare			0.575%			1.621%			-0.464%

Table 6: The Effects of Transaction Costs. This table shows the results obtained from simulating the model setting λ equal to 6, 8, and 10 percent. $\overline{\Delta c_t}$ and $\overline{\Delta h_t}$ are average non-durable consumption and house consumption growth. $\sigma(\Delta c_t)$ and $\sigma(\Delta h_t)$ are the respective standard deviations. $\overline{H_t}$ is average house size. $\overline{P_t H_t}$ is average house value. $\overline{Fin.Assets}$ is average financial savings. The table also shows the total number of house trades, as well as the number of house trades associated with an increase ($H_t > H_{t-1}$) and a decrease ($H_t < H_{t-1}$) in house size. $\sigma(SDF)$ is the volatility of the stochastic discount factor. Welfare measures the percentage increase in consumption in the benchmark case needed to ensure the same level of lifetime utility. The table shows results conditional on the level of interest rates and unconditional values.

	$\lambda = 0.06$			$\lambda = 0.08$			$\lambda = 0.10$		
	r_t low	r_t high	Uncond.	r_t low	r_t high	Uncond.	r_t low	r_t high	Uncond.
$\overline{\Delta c_t}$	0.095	-0.001	0.047	0.100	-0.010	0.045	0.097	-0.009	0.044
$\sigma(\Delta c_t)$	0.102	0.134	0.128	0.115	0.147	0.142	0.114	0.146	0.141
$\overline{\Delta h_t}$	0.175	-0.087	0.044	0.083	0.018	0.051	0.076	0.025	0.051
$\sigma(\Delta h_t)$	0.386	0.414	0.421	0.285	0.201	0.249	0.278	0.161	0.229
$\overline{H_t}$	73.553	35.656	54.678	54.300	44.719	49.528	51.155	43.476	47.330
$\overline{P_t H_t}$	100.885	36.843	68.988	71.247	48.189	59.763	67.204	46.978	57.130
$\overline{Fin.Assets}$	9.633	4.853	7.252	10.122	4.438	7.291	10.043	4.463	7.263
# House Trades	3396	2293	5689	1576	598	2174	1386	501	1887
$\#H_t > H_{t-1}$	3396	573	3969	1576	482	2058	1386	493	1882
$\#H_t < H_{t-1}$	0	1720	1720	0	116	116	0	5	5
$\sigma(SDF)$	0.184	0.4679	0.394	0.206	0.461	0.394	0.207	0.449	0.385
Welfare			0.337%						-0.412%

Table 7: The Effects of Risk Aversion. This table shows the results obtained from simulating the model setting γ equal to 3 and 5. $\overline{\Delta c_t}$ and $\overline{\Delta h_t}$ are average non-durable consumption and house consumption growth. $\sigma(\Delta c_t)$ and $\sigma(\Delta h_t)$ are the respective standard deviations. $\overline{H_t}$ is average house size. $\overline{P_t H_t}$ is average house value. $\overline{Fin.Assets}$ is average financial savings. The table also shows the total number of house trades, as well as the number of house trades associated with an increase ($H_t > H_{t-1}$) and a decrease ($H_t < H_{t-1}$) in house size. $\sigma(SDF)$ is the volatility of the stochastic discount factor. The table shows results conditional on the level of interest rates and unconditional values.

	$\gamma = 3$			$\gamma = 5$		
	r_t low	r_t high	Uncond.	r_t low	r_t high	Uncond.
$\overline{\Delta c_t}$	0.100	-0.010	0.045	0.085	-0.000	0.042
$\sigma(\Delta c_t)$	0.115	0.147	0.142	0.094	0.130	0.121
$\overline{\Delta h_t}$	0.083	0.018	0.051	0.071	0.024	0.047
$\sigma(\Delta h_t)$	0.285	0.201	0.249	0.247	0.1473	0.205
$\overline{H_t}$	54.300	44.719	49.528	41.758	36.108	38.944
$\overline{P_t H_t}$	71.247	48.189	59.763	54.634	38.641	46.669
$\overline{Fin.Assets}$	10.122	4.438	7.291	9.519	4.674	7.106
# House Trades	1576	598	2174	1535	566	2101
$\#H_t > H_{t-1}$	1576	482	2058	1535	553	2088
$\#H_t < H_{t-1}$	0	116	116	0	13	13
$\sigma(SDF)$	0.206	0.461	0.394	0.250	0.694	0.570

Table 8: Hedging Demands. This table shows the results from simulating the model setting $\delta_t = 0.10$ for $t < 34$ and $\delta_t = 0.15$ for $t \geq 34$ for different values of initial wealth, W_0, and the coefficient of relative risk aversion γ. $\overline{H_t}$ is average house size. # Trades is the cumulative number of house trades. The table also shows the percentage difference in average house size and the number of house trades relative to the $\delta_t = 0.10\ \forall t$ scenario.

	Age					
	25	27	29	31	33	34
	$\gamma = 3.00, W_0 = 0.00$					
$\overline{H_t}, \delta_t = 0.15$ for $t \geq 34$	22.151	35.763	44.216	51.011	56.136	62.356
$\overline{H_t}, \delta_t = 0.10\ \forall t$	24.030	35.990	43.103	47.825	50.993	52.925
% Dif. $\overline{H_t}$	-7.818	-0.632	2.584	6.662	10.087	17.820
# Trades, $\delta_t = 0.15$ for $t \geq 34$	-	917	1396	1741	1965	2212
# Trades, $\delta_t = 0.10\ \forall t$	-	836	1257	1520	1690	1777
% Dif. # Trades	-	9.689	11.058	14.539	16.272	24.479
	$\gamma = 3.00, W_0 = 10.00$					
$\overline{H_t}, \delta_t = 0.15$ for $t \geq 34$	46.530	51.000	55.111	59.090	62.205	67.143
$\overline{H_t}, \delta_t = 0.10\ \forall t$	45.278	50.969	54.205	57.065	58.841	60.195
% Dif. $\overline{H_t}$	2.766	0.060	1.672	3.549	5.717	11.542
# Trades, $\delta_t = 0.15$ for $t \geq 34$	-	281	500	662	779	963
# Trades, $\delta_t = 0.10\ \forall t$	-	288	505	669	770	839
% Dif. # Trades	-	-2.431	-0.990	-1.046	1.169	14.779
	$\gamma = 5.00, W_0 = 0.00$					
$\overline{H_t}, \delta_t = 0.15$ for $t \geq 34$	19.611	27.736	34.040	39.328	43.989	49.878
$\overline{H_t}, \delta_t = 0.10\ \forall t$	19.611	27.601	32.841	36.601	39.786	41.635
% Dif. $\overline{H_t}$	0.000	0.489	3.650	7.449	10.563	19.797
# Trades, $\delta_t = 0.15$ for $t \geq 34$	-	778	1253	1606	1873	2167
# Trades, $\delta_t = 0.10\ \forall t$	-	778	1207	1469	1683	1794
% Dif. # Trades	-	0.000	3.811	9.326	11.289	20.792

Figure 1: Income Profile

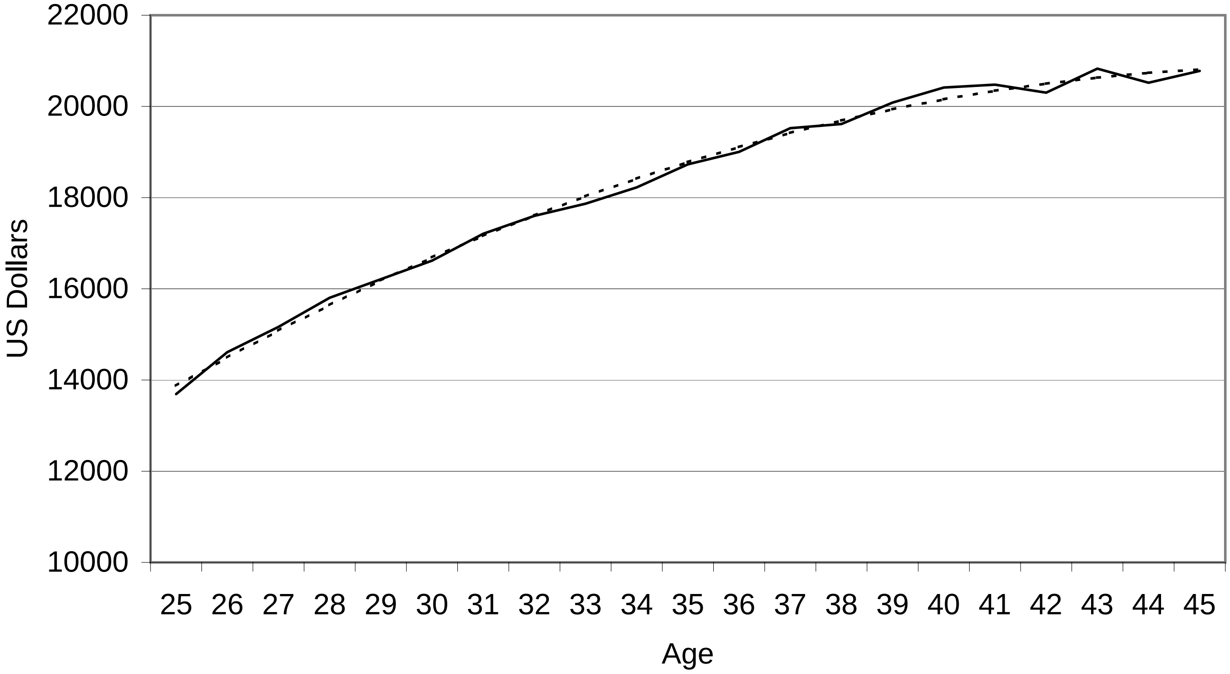

Figure 2: Logarithm of Real House Prices in the US

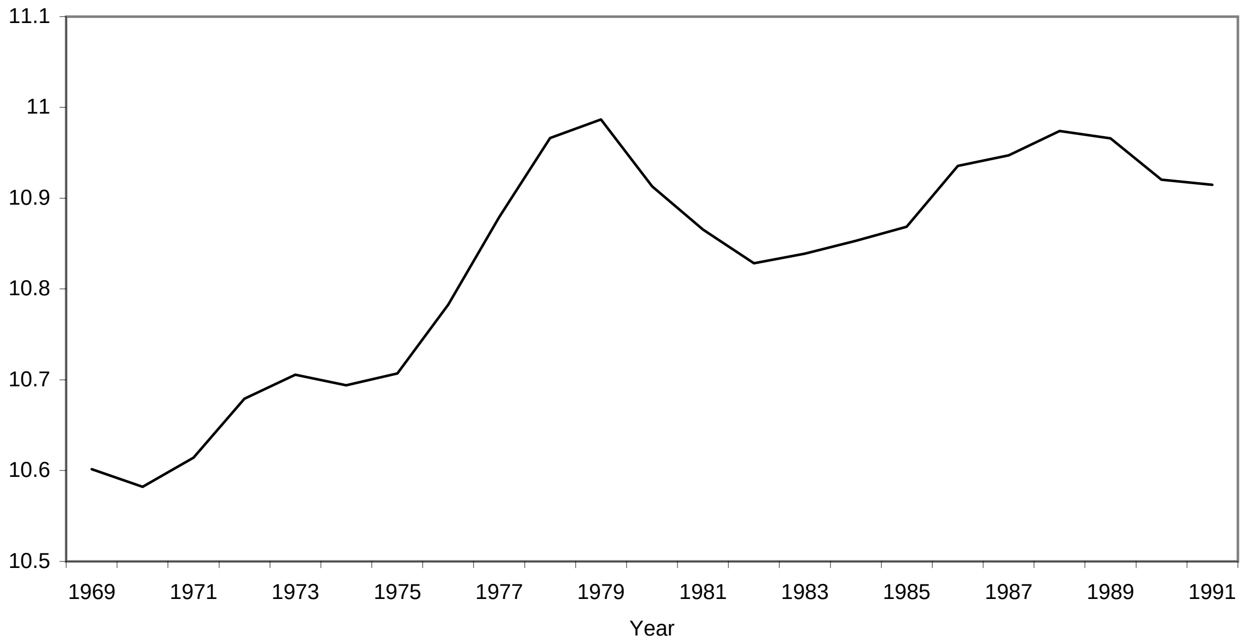

Figure 3: Income Shocks, House Price Shocks and Interest Rates

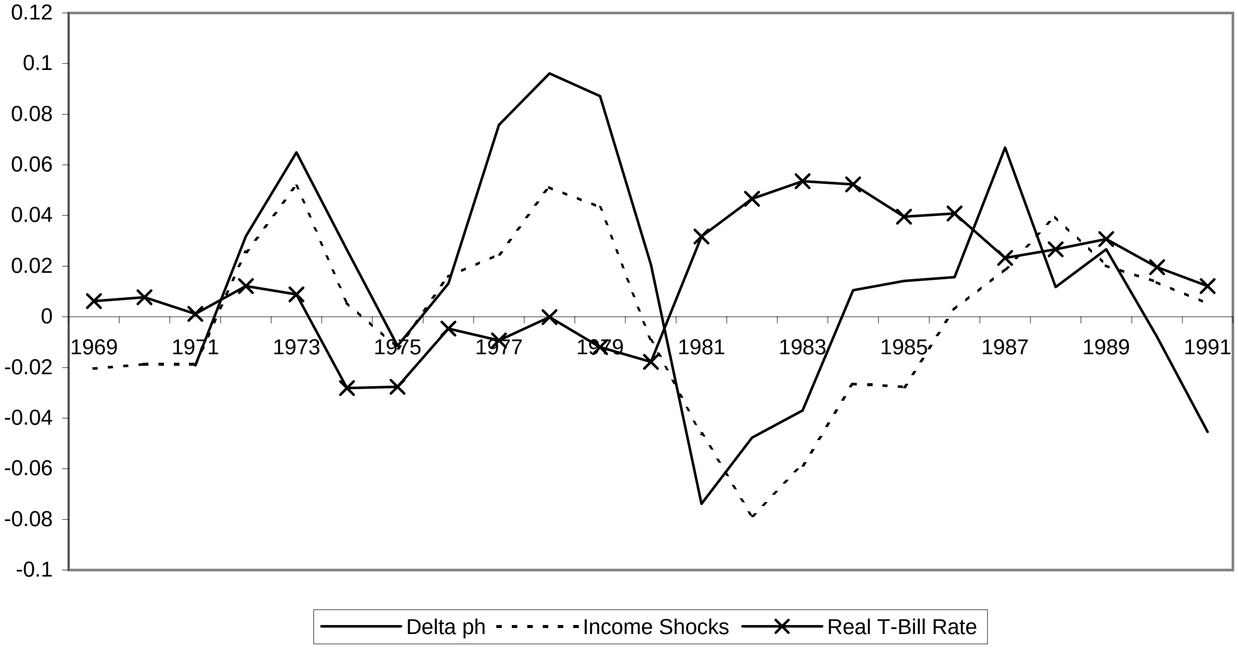

Figure 4 - Benchmark Results

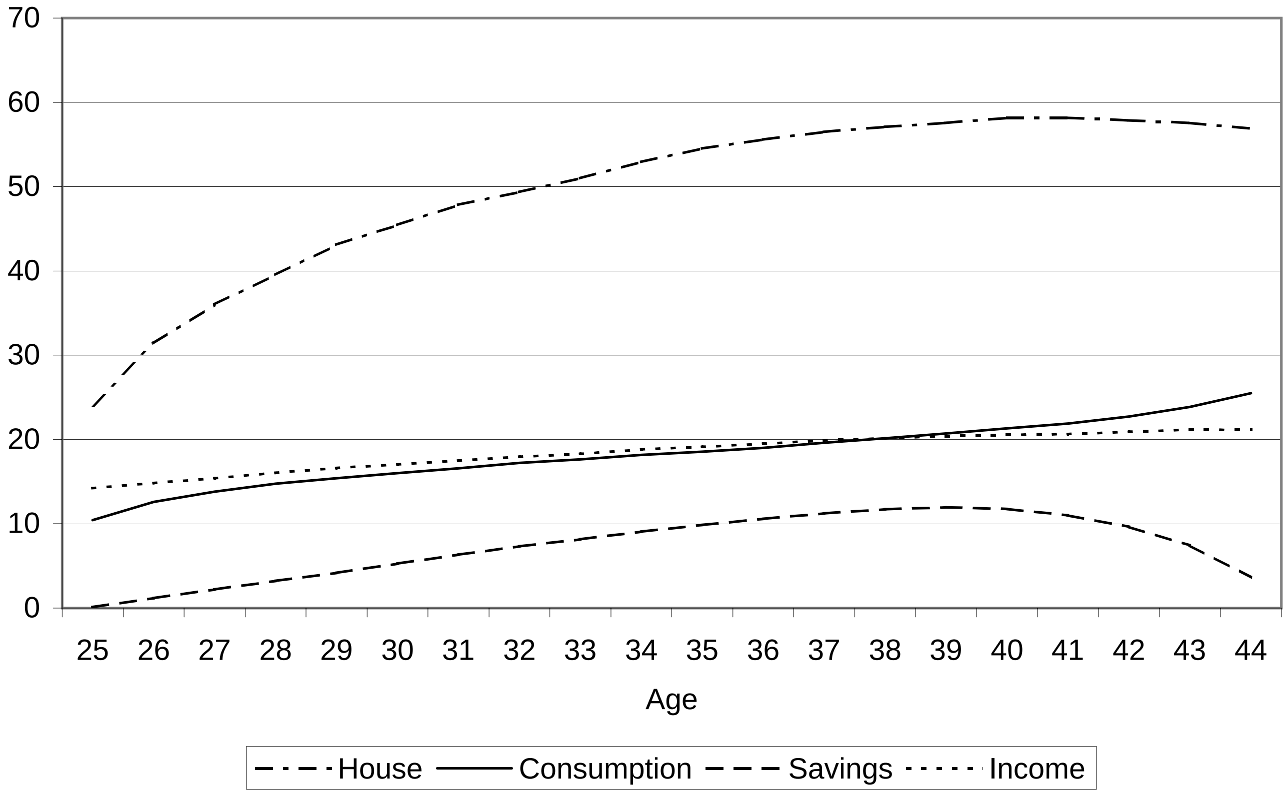

Figure 5 - Hedging Demands