Housing, Consumption, and Asset Pricing*

Monika Piazzesi, Martin Schneider, and Selale Tuzel

May 2003

Abstract

This paper builds an equilibrium asset pricing model with housing consumption. Agents care about the composition of a consumption basket that contains shelter and other goods. The presence of *composition risk* increases the mean and variance of excess stock returns and lowers the riskfree rate. Stock prices exhibit mean reversion because they depend on the expenditure share of non-housing consumption. We show that this state variable indeed forecasts excess returns. It also helps account for the value and small firm premia.

*This draft is still preliminary and incomplete. UCLA Anderson School and Economics Department, Los Angeles CA 90095. Email: piazzesi@ucla.edu, martin@econ.ucla.edu and selale.tuzel@anderson.ucla.edu. We would like to thank John Cochrane, Marjorie Flavin, Lars Hansen, Per Krusell, Martin Lettau, Hanno Lustig, Sydney Ludvigson, Stijn Van Nieuwerburgh, Pedro Santa-Clara, Steve Ross and seminar participants at Caltech, Kellogg, UC San Diego, NYU, Princeton, Texas Austin, UC Berkeley, and USC. The latest version of this paper can be downloaded from http://www.econ.ucla.edu/schneider.

1 Introduction

Recent work in asset pricing has tried to construct a macroeconomic model in which excess stock returns are (i) on average much higher than both the riskless rate and excess bond returns, (ii) predictable, and (iii) much more volatile than the riskless rate and excess bond returns.[1] The standard consumption-based (CCAPM) approach assumes that agents consume a single numeraire good. With power utility, the model predicts that excess stock returns are low, unpredictable and not very volatile.

This paper proposes a simple modification of the CCAPM: agents consume both *housing services*, s_t, and a numeraire (non-housing) consumption good, c_t. We calibrate a version with power utility over a CES aggregator of housing and non-housing consumption to NIPA data. A key feature of our model is that agents care about the composition of their consumption basket. Fluctuations in expenditure shares are a source of risk. The presence of this *composition risk* leads the model to do a better job on properties $(i)-(iii)$ above than the standard CCAPM.

First, we show that stocks expose their owners to composition risk and thus command a higher equity premium. At the same time, composition risk strengthens incentives for precautionary saving, which lowers the riskfree rate. Second, asset prices depend on state variables that capture composition, which are typically persistent, but mean-reverting. In particular, the model predicts that, in addition to the dividend yield, the *expenditure share on non-housing consumption*, α_t, should forecast excess stock returns. We document that this is indeed the case in the data. Third, changes in composition induce price volatility. As in the data, this volatility is not due to news about cash flow; in fact, dividends are not forecastable in our model. Finally, we explore the role of composition risk for explaining value and small firm premia. Our model motivates a two factor regression model with the expenditure share α_t as a scaling variable. This regression accounts for more than three quarters of the variation in returns on 25 size and book-to-market portfolios.

Composition risk is priced because, with nonseparable utility, the marginal rate of substitution depends not only on numeraire consumption growth, but also on the change in the relative quantities of numeraire and housing services:

$$M_{t+1} = \beta \frac{u_1\left(c_{t+1}, s_{t+1}\right)}{u_1\left(c_t, s_t\right)} = \beta \left(\frac{c_{t+1}}{c_t}\right)^{-\frac{1}{\sigma}} \left(\frac{g_1\left(1, s_{t+1}/c_{t+1}\right)}{g_1\left(1, s_t/c_t\right)}\right)^{\frac{\sigma-\varepsilon}{\sigma}}, \tag{1}$$

where β is the discount factor, σ and ε are the inter- and intratemporal elasticities of substitution, respectively, and g is a CES aggregator. As usual, agents care about consumption risk: numeraire is valued more in states of the world where consumption falls. The effect of composition risk depends on whether agents desire to smooth consumption over time more or less than across goods. For example, suppose *inter*temporal smoothing is more important $(\sigma < \varepsilon)$. Then numeraire is less desirable in states of the world where the housing stock already yields a relatively high level of services $(u_{12} < 0)$. Agents thus

[1] See Cochrane (1997) and Campbell (2001) for surveys of recent literature.

prefer assets that insure them against states where there is a relative shortfall of housing. The pricing kernel (1) weights payoffs more in states where s/c falls, because it means that g_1 falls[2] and its exponent is negative. In contrast, if intratemporal smoothing is more important ($\sigma > \varepsilon$), the exponent is positive. Agents now prefer assets that pay off when numeraire consumption drops relative to that of shelter, which means that s/c increases.

To measure composition risk, we focus on the expenditure share of non-housing consumption. In our model, the change in α_t may be viewed as a second "risk factor", in addition to numeraire consumption growth. Indeed, with a perfect rental market,

$$M_{t+1} = \beta \left(\frac{c_{t+1}}{c_t} \right)^{-\frac{1}{\sigma}} \left(\frac{\alpha_{t+1}}{\alpha_t} \right)^{\frac{\varepsilon - \sigma}{\sigma(\varepsilon - 1)}} . \tag{2}$$

We say that there is a *positive composition effect* if numeraire is desired in states where its expenditure share has increased. The composition effect is thus positive, if the exponent on expenditure growth α_{t+1}/α_t is positive. If intertemporal smoothing is important ($\varepsilon > \sigma$), a positive composition effect obtains if the intratemporal elasticity of substitution is not too small ($\varepsilon > 1$). Agents then prefer assets that insure them against states with higher non-housing expenditures.[3] We provide empirical evidence about the intratemporal elasticity of substitution ε using aggregate NIPA data. We find a value for ε above but close to 1, while conventional estimates of the intertemporal elasticity σ are typically much lower than 1. This evidence is thus consistent with a positive composition effect.

In the data, our numeraire (non-housing) consumption growth series behaves much like aggregate consumption growth: it is smooth, close to serially uncorrelated, almost homoskedastic, and not very correlated with stock returns (denominated in units of numeraire). These are the properties that doom the standard model; here they imply that our first risk factor, numeraire consumption growth, cannot matter much. To give an idea about the behavior of our second risk factor, Figure 1 plots the non-housing expenditure share: it is both persistent and heteroskedastic, and its overall correlation with the dividend yield is positive. Historically, the expenditure share on non-housing consumption has fallen during stock market booms. It also was much more volatile during the 1940s, when the expenditure share was large.

Average excess returns for an asset with return R_{t+1} are approximately equal to $-cov\,(R_{t+1}, M_{t+1})$. Figure 1 correctly suggests that, in the post-depression era, the change in the non-housing share α_{t+1}/α_t is negatively correlated with stock returns. With a positive composition effect, composition risk thus increases the equity premium. At the same time, composition risk *lowers* the average riskless rate, that is, it lowers $1/E\,[M_{t+1}]$. Precautionary savings for composition risk pushes the interest rate down, and thus helps avoid the riskfree rate puzzle.

[2] The partial g_{12} is positive for any CES aggregator g.

[3] A positive composition effect also obtains if numeraire and housing are highly complementary, i.e. $\varepsilon < \min\{\sigma, 1\}$.

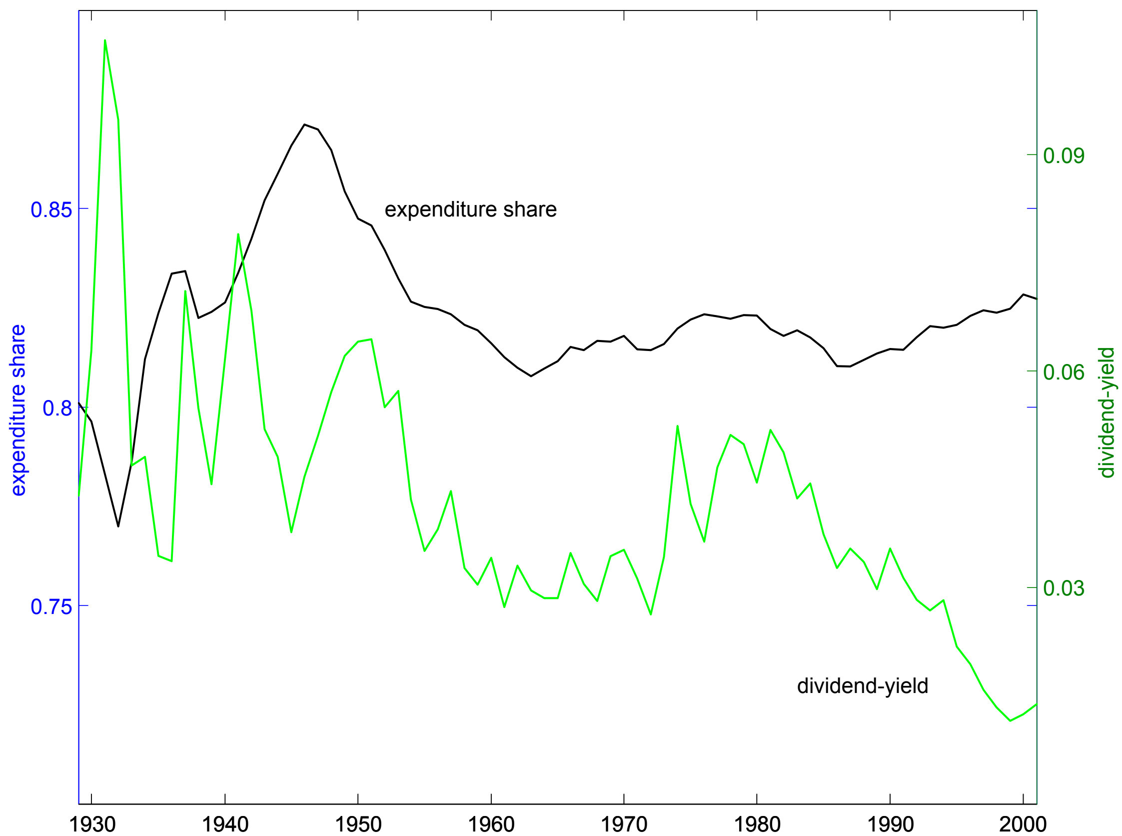

Figure 1: Annual NIPA data on expenditures and Shiller data on dividend-yields.

Figure 1 shows that the non-housing expenditure share α_t changes slowly and reverts to its mean. In our model, these properties carry over to interest rates. For example, suppose that α_t is high, which means that agents consume more non-housing items than housing. In other words, there is relative housing shortfall. Agents will expect the housing stock to grow and thus expect to be able to substitute housing for consumption. Because of consumption smoothing, agents will borrow, which increases the riskfree rate and longer yields. Interest rates thus comove with α_t. But when α_t is high, it is also volatile. This strengthens incentives for precautionary savings. This effect dampens fluctuations in the riskless rate, so that its volatility is not excessive and bond premia are not high or volatile.

In our model, composition risk also leads to slow mean reversion in asset prices. Again, suppose that α_t is high, so that agents expect the housing stock to grow. Assets that pay off in numeraire, such as bonds and stocks, are then less desirable and trade at lower prices relative to their dividends. At the same time, agents know that, once more houses are built, price-dividend ratios must increase. The model thus predicts that both the dividend yield and the non-housing share should forecast excess stock returns. Moreover, the explanatory power of a regression of excess returns on either of these variables should increase with the forecast horizon. In light of Figure 1, it is not surprising that this is indeed true in the data. For example, high ex post risk premia observed in the mid 1970s coincided with an increase in the non-housing share to levels last seen during the war. Our model suggests that this reflects high *ex ante* risk premia due to composition risk. As in the data, excess holding returns on bonds are also predictable, with expected premia moving together with expected excess stock returns. However, since bonds do not promise dividend payoffs that grow over time, their sensitivity to changes in discount rates due to composition risk is significantly lower than that of stocks.

We also explore the role of composition risk for small firm and value premia. An interesting property of our pricing kernel is that the current level of the expenditure share α_t matters for how risk premia depend on the correlation between returns on the one hand and innovations to consumption growth and expenditure growth on the other. This observation motivates an empirical two-factor model with consumption growth and the growth rate of α_t as factors and α_t itself as a scaling variable. We examine this model using the Fama-MacBeth regression methodology. We find that it gives rise to large and significant pricing errors. In addition, the regression coefficients are not easily reconciled with the parameter values used for our time series results.

Nevertheless, it is interesting that the R^2 in the cross sectional regression of mean returns on betas is over 80%, and nearly equals that found for the Fama-French three factor model. This is due to several effects. First, the CCAPM itself already does quite well in the post-war era. Spectacular failures of the CCAPM appear to arise only if the sample is chosen to either include the Great Depression or exclude the 1940s and 1950s. Second, the growth rate of α_t is a significant factor. Third, performance increases further if either factor is scaled by α_t. This result is reminiscent of that in Lettau and Ludvigson (2001). Scaling consumption growth by a variable that is high in stock market slumps (in their case, the consumption-wealth ratio) increases the R^2, because small value stocks are more correlated with consumption growth in bad times. Figure 1 suggests that the expenditure share α_t is such a variable.

Our model also has implications for the value of the aggregate housing stock, the claim to future housing services. In the model, returns on housing are somewhat smaller on average and less volatile than stock returns. This is in line with a corresponding measure of housing returns we construct from the data. At the same time, our model has little to say about house price risk experienced by individuals. As documented by Flavin and Yamashita (2002), aggregate risk is only about one quarter of housing return volatility at the individual level. Our analysis effectively assumes that the remainder of this risk can be diversified away. Even though this may not be literally true, we believe that a representative agent model is a natural starting point for an analysis of consumption-based asset pricing with housing. For one thing, the composition risk effect we emphasize is likely to be present in any model, even if agents are heterogeneous. In addition, our results suggest that even though our model cannot account for *individual* house price volatility, it can go some way to getting at aggregate volatility.

To match historical risk premia quantitatively, an asset pricing model must either imply a high market price of risk or a high quantity of risk. The CCAPM implies neither, since both the risk aversion parameter $1/\sigma$ and the volatility of consumption growth are low. Our model can match average premia if the composition risk term is volatile enough. It is intuitive to define the quantity of composition risk as the volatility of relative quantities s/c. The link between the quantity of risk and the premium then depends on the elasticities of substitution. For a given expenditure share process α, the premium is high if the exponent on α_{t+1}/α_t in (2) is high. At the same time, for given α, the elasticity ε of intratemporal substitution governs the volatility of relative quantities. With low risk aversion, values of ε that achieve a realistic equity premium also imply an implausibly high volatility of relative quantities. As ε is increased to reduce the quantity of composition risk, the model-implied equity premium falls. This suggests that consumption and composition risk are not the only sources of risk in the data. Nevertheless, our results suggest that at least some of the premium is due to composition risk. Exactly how much of the premium is due to composition risk can only be determined with reliable quantity data. This is an important issue for future research.

The paper proceeds as follows. Section 2 discusses related work. Section 3 presents the model and derives our pricing equations. Section 4 documents data properties and specifies the forcing process for the equilibrium model. Section 5 discusses evidence on utility parameters. Section 5.2 investigates the Euler equation using data on returns. Section 6 documents properties of equilibrium returns. Section 7 reports time-series evidence, while Section 8 reports cross-sectional asset pricing evidence.

2 Related Work

To our knowledge, there is no prior work that examines the effects of housing on asset prices in the context of a general equilibrium model. The paper most closely related to ours is by Flavin (2001), who examines the portfolio choice problem of an agent who invests in both financial assets and real estate. Her setup allows for both adjustment costs to housing and nonseparabilities in the utility function. Working in continuous time, she shows that if returns on houses and stocks are uncorrelated, as is the case at the individual level, the CAPM must hold, while nonseparability prevents the standard CCAPM from holding. The latter property is also crucial for our analysis. However, in our paper asset prices are derived endogenously and their properties are explored quantitatively. Other papers that consider portfolio choice with exogenous returns include Flavin and Yamashita (2002), and Cocco (2000). Examples of general equilibrium models of housing are Heathcote and Morris (2001), who explore the business cycle implications of an RBC model with a construction sector, and Ortalo-Magne and Rady (1999), who analyze an overlapping generations model to study prices and volume in the housing market. Neither of these papers is concerned with financial assets.

Dunn and Singleton (1986), Eichenbaum and Hansen (1990) and Heaton (1993, 1995) explore asset pricing with durable goods. They find that distinguishing between durable goods on the one hand and nondurables and services on the other does not matter much for asset pricing. While these papers carefully model the service flows from durable good purchases, they lump services from housing together with other nondurables and services. However, housing is the most durable good, because it consists of structures and land, and land provides services flows forever. If durability matters for asset pricing, we are thus more likely to find its effect by separating housing from other durables. Eichenbaum, Hansen and Singleton (1988), Jagannathan and Wang (1996), and Santos and Veronesi (2001) distinguish between consumption and leisure. The latter paper predicts (and verifies) that the ratio of consumption to labor income forecasts stock returns. However, the pricing kernel is the same as in the standard model. Hence, the result does not arise from composition risk as we have defined it; instead, it obtains because the riskiness of stocks changes over time with the correlation of dividends and aggregate consumption.

To date, the most successful representative agent asset pricing model is due to Campbell and Cochrane (1999). They propose a model in which agents consume a single good, but want to "catch up with the Joneses." Their pricing kernel also contains a persistent, heteroskedastic state variable, which they call the consumption-surplus ratio. However, a key difference is that the non-housing share α_t is observable, while the consumption-surplus ratio is a parametric function of past aggregate consumption, with parameters that must be inferred from asset market data.[4] While our model does not perform as well as the Campbell-Cochrane model on the issues $(i) - (iii)$ above, its pricing kernel is arguably more closely tied to macro data. Moreover, we use moderate values for the coefficient of relative risk aversion in our calibration.

The success of the Fama-French three-factor-model and the failure of the CCAPM in pricing the cross-section of stock returns has spawned a recent literature that looks for macroeconomic factors other than consumption growth. These papers employ the Fama-French regression framework, but motivate the choice of regressors with economic arguments. For example, Jagannathan and Wang (1996) argue that returns on a broad stock index do not adequately capture the return on wealth, since they do not reflect human capital. This motivates using labor income growth as an additional factor. For similar reasons, Kullmann (2002) includes returns on real estate as a factor.

Lettau and Ludvigson (2001a,b) suggest that, if the CCAPM holds conditionally, but not unconditionally, better performance may obtain if the CCAPM is scaled by variables that capture relevant conditioning information. Based on the representative agent's budget constraint, they propose cay, an estimate of the consumption-wealth ratio, as a scaling variable. This variable forecasts returns and improves pricing in the cross-section. Lustig and Nieuwerburgh (2002) postulate a linear factor model where an estimate of the housing to human wealth ratio, which they call my, is a scaling variable. To motivate this formulation, they sketch, but do not solve, an extension of our model that allows for heterogenous agents and collateral constraints, adapting the setup of Lustig (2001).

[4] A second difference is that the share α_t is bounded. With a positive composition effect, the second term in marginal utility, $\alpha_t^{\frac{\varepsilon-\sigma}{\sigma(\varepsilon-1)}}$, is thus bounded above by one, and marginal utility itself is bounded above by the standard expression $c_t^{-\gamma}$. This is in contrast to the Campbell-Cochrane model, where marginal utility increases without bound as the consumption-surplus ratio goes to zero. Our results thus do not rely on having very high risk aversion in certain parts of the state space.

Both our time series and cross sectional results confirm the findings in Cochrane (1991, 1996) who investigates real estate investment as pricing factor in a production-based approach. Cochrane (1991) documents that real-estate investment growth predicts stock returns. Cochrane (1996) finds that real-estate investment growth matters for the cross section of stock returns. Moreover, the important component in real estate investment is residential real estate, not commercial real estate investment (Table 9 on page 615). This finding supports our approach of introducing real estate using a consumption-based view, where residential real estate matters to consumers.

3 Model

3.1 Setup

A large number of identical agents enjoy both housing services h_t and a nondurable, numeraire consumption good c_t. Preferences are

$$E\left[\sum_{t=0}^{\infty}\beta^t u\left(c_t, s_t\right)\right], \tag{3}$$

where

$$\begin{aligned} u\left(c_t, s_t\right) &= \frac{g\left(c_t, s_t\right)^{1-\frac{1}{\sigma}}}{1-\frac{1}{\sigma}}, \\ g\left(c_t, s_t\right) &= \left(c_t^{\frac{\varepsilon-1}{\varepsilon}} + \omega s_t^{\frac{\varepsilon-1}{\varepsilon}}\right)^{\frac{\varepsilon}{\varepsilon-1}}, \end{aligned}$$

for positive scalars σ and ε. The CES aggregator g depends on the importance of intratemporal smoothing. The higher the intratemporal elasticity of substitution ε, the more willing the agents are to substitute housing and numeraire consumption within the same period. The two goods are perfect substitutes if ε is infinite, while they are perfect complements if ε is zero.[5] Taking the limit as $\varepsilon \to 1$ yields a Cobb-Douglas aggregator. The period utility function u depends on how willing agents are to substitute the CES aggregated bundle over time. The higher the intertemporal elasticity of substitution σ, the more willing the agents are to substitute. The intratemporal elasticity is also the inverse of the coefficient of relative risk aversion, $\gamma = 1/\sigma$.

There are two real assets in the economy. A single Lucas tree produces a stream d_t of numeraire. Shares of this tree trade at the price p_t^s. Agents also accumulate housing

[5] We use standard Hicksian language here: two goods are substitutes if and only if $\varepsilon > 1$. This property can be inferred from data on relative prices and quantities, and has nothing to do with the agent's intertemporal concern for smoothing consumption. Some papers refer to $u_{12} < 0$ as the case where numeraire and shelter are "substitutes", while the case $u_{12} > 0$ is referred as "complements". We refrain from this language here, since the second derivative of the utility function captures both intertemporal and intratemporal tradeoffs.

capital. One unit of housing capital installed in period t produces η units of housing services in period $t+1$.[6] Residential investment is perfectly reversible and divisible. Housing capital trades at the price p_t^h and depreciates at the rate δ. A consumer's budget constraint is

$$c_t + p_t^h h_t + p_t^s \theta_t = (p_t^s + d_t)\,\theta_{t-1} + p_t^h h_{t-1}(1-\delta) + p_t^h \overline{h}_t^n, \tag{4}$$

where θ_t denotes holdings of the tree and $\overline{h}_t^n = \overline{h}_t - (1-\delta)\,\overline{h}_t - 1$ denotes new housing construction.

The economy is summarized by the preference parameters β, ω, σ and ε, as well as a stochastic process for output and the total supply of housing capital, $\left(y_t, \overline{h}_t\right)$. In this economy, output equals dividends $y_t = d_t$. An equilibrium is a collection of processes $\left(c_t, s_t, h_t, \theta_t, p_t^s, p_t^h\right)$ such that (i) (c_t, s_t, h_t) solve the consumer's problem of maximizing utility (3) subject to the budget constraint (4) as well as the technical constraint $s_t = \eta h_{t-1}$, and (ii) markets clear for stocks $(\theta_t = 1)$, housing capital $\left(h_t = \overline{h}_t\right)$ and consumption $(c_t = y_t)$.

3.2 Bond, Stock, and House Prices

The Lagrangian for the consumer's problem is

$$L = E\left[\sum_{t=0}^{\infty} \beta^t \left(u\left(c_t, \eta h_{t-1}\right) - \mu_t \left[c_t + p_t^h h_t + p_t^s \theta_t - (p_t^s + d_t)\,\theta_{t-1} - p_t^h h_{t-1}(1-\delta) - p_t^h \overline{h}_t^n\right]\right)\right]$$

where μ_t is the multiplier on the budget constraint. The first order conditions are

$$\begin{aligned}
\frac{\partial L}{\partial c_t} &= \beta^t \left(u_1\left(c_t, s_t\right) - \mu_t\right) = 0 \\
\frac{\partial L}{\partial h_t} &= \beta^t \left(-\mu_t p_t^h + \beta E_t\left[\mu_{t+1}(1-\delta)\, p_{t+1}^h + \eta u_2\left(c_{t+1}, s_{t+1}\right)\right]\right) = 0 \\
\frac{\partial L}{\partial \theta_t} &= \beta^t \left(-\mu_t p_t^s + \beta E_t\left[\mu_{t+1}\left(p_{t+1}^s + d_{t+1}\right)\right]\right) = 0
\end{aligned}$$

These conditions can be rewritten using the pricing kernel, which represents the present value of a unit of numeraire one period ahead:

$$M_{t+1} = \beta \frac{u_1\left(c_{t+1}, s_{t+1}\right)}{u_1\left(c_t, s_t\right)} = \beta \left(\frac{c_{t+1}}{c_t}\right)^{-\frac{1}{\sigma}} \left(\frac{1+\omega\left(\frac{s_{t+1}}{c_{t+1}}\right)^{\frac{\varepsilon-1}{\varepsilon}}}{1+\omega\left(\frac{s_t}{c_t}\right)^{\frac{\varepsilon-1}{\varepsilon}}}\right)^{\frac{\sigma-\varepsilon}{\sigma(\varepsilon-1)}}$$

The first term is familiar from the standard power utility model. Numeraire is valued more highly tomorrow when consumption is lower than today. The more willing the

[6] This timing convention is innocuous. All stock variables represent end-of-period values. For example, the price of a stock measures the value a dividend stream y that starts in $t+1$.

consumer is to substitute over time (higher σ), the smaller is the effect of consumption growth on valuation. The second term captures the adjustment in the value of numeraire that is required by the presence of housing.

This adjustment depends on consumers' attitudes towards changes in the *composition* of consumption bundles over time. If the intertemporal elasticity of substitution is smaller than the intratemporal elasticity ($\sigma < \varepsilon$, or, equivalently, $u_{12} < 0$), the consumer is more concerned about smoothing the level of utility than about keeping the composition of bundles similar. In particular, if there is a shortfall of shelter, he is eager to smooth utility by substituting numeraire. Consequently, an asset denominated in numeraire is more attractive if it pays out in states of the world where there is a relative shortfall of shelter. In contrast, if $\sigma > \varepsilon$, the consumer will try to smooth relative quantities over time, and an asset is more attractive if it pays out in states where the relative quantity of numeraire is lower than today. The presence of housing reduces the need for numeraire in states of relatively high numeraire growth: it is difficult to make up for housing with other goods (low ε), and this smoothing is not very important in the first place (high σ).

Using the pricing kernel notation, the first order conditions become

$$p_t^s = E\left[M_{t+1}\left(p_{t+1}^s + d_{t+1}\right)\right] \tag{5}$$

$$p_t^h = E\left[M_{t+1}\left(p_{t+1}^h\left(1-\delta\right) + \eta\frac{u_2\left(c_{t+1}, s_{t+1}\right)}{u_1\left(c_{t+1}, s_{t+1}\right)}\right)\right] \tag{6}$$

As usual, the value p^s of the tree is the present value of future dividends d, discounted by M_{t+1}. More generally, any financial asset can be priced in this way. The value p^b of a consol bond is the present discounted value of dividends $d = 1$. The same is true for the value p^h of housing capital, if the dividend is taken to be utility services provided in the following period.

3.3 Asset Prices and Expenditure Shares

The presence of composition risk implies a simple relationship between asset prices and expenditure shares. Suppose there is a rental market for housing, so that housing services can be bought at the price q_t. Consumers thus minimize expenditure $c_t + q_t s_t$ to achieve a desired utility level, $u\left(c_t, s_t\right) = \overline{u}$, say. The first order condition for this problem imply

$$q_t = \frac{u_2\left(c_t, s_t\right)}{u_1\left(c_t, s_t\right)} = \omega\left(\frac{s_t}{c_t}\right)^{-\frac{1}{\varepsilon}}. \tag{7}$$

We will use this relationship between relative prices and quantities in Section 5 to obtain an estimate of ε. It follows that the share of non-housing expenditure in total expenditure is

$$\alpha_t := \frac{c_t}{c_t + q_t s_t} = \frac{1}{1+\omega\left(\frac{s_t}{c_t}\right)^{\frac{\varepsilon-1}{\varepsilon}}}. \tag{8}$$

The pricing kernel can now be written as in equation (1).

Consumers' attitudes towards changes in expenditure shares α depend not only on their attitude towards the composition of consumption bundles, but also on the link between relative quantities and expenditure shares. By (8), relative quantities move together with expenditure shares if and only if the intratemporal elasticity of substitution ε is larger than one, that is, when housing and non-housing consumption are substitutes. In this case, an increase in the relative price of numeraire, say, leads to a sufficiently large decrease in the relative quantity consumed that the expenditure share also falls. It also follows that prices and shares α move in opposite directions iff $\varepsilon > 1$:

$$q_t = \omega^{\frac{\varepsilon}{\varepsilon-1}} \left(\frac{\alpha_t}{1-\alpha_t} \right)^{\frac{1}{\varepsilon-1}}. \tag{9}$$

How the valuation of numeraire moves with changes in the expenditure ratio depends on the intratemporal and intertemporal elasticities. We say that there is a *positive composition effect*, if numeraire is more valued in states in which its expenditure share is higher than today (that is, $(\varepsilon - \sigma)/\sigma(\varepsilon - 1) > 0$).[7] There are two sets of parameters that give rise to a positive composition effect. First, suppose that the intratemporal elasticity ε is higher than the intratemporal elasticity ε and that the goods are substitutes. With $\varepsilon > \sigma$ consumers will want to substitute numeraire to make up for a shortfall in housing. In addition, with $\varepsilon > 1$ a shortfall in housing precisely coincides with an increase in the non-housing expenditure share. The second case occurs if $\sigma > \varepsilon$ and the two goods are complements. If $\sigma > \varepsilon$, consumers try to smooth relative consumption over time. With $\varepsilon < 1$, expenditure shares and relative quantities move opposite to each other, so numeraire is again needed when the expenditure share of numeraire is high.

3.4 Comparison to the CCAPM

Further intuition about the pricing properties of the model can be gained by considering an approximation to the stochastic discount factor. We denote by z_t the consumption-housing expenditure ratio,

$$z_t = \frac{\alpha_t}{1-\alpha_t} = \frac{c_t}{q_t s_t}.$$

The ratio z_t is increasing in α_t, but ranges over the whole real line. We now write $M_{t+1} = \beta e^{-(1/\sigma)\Delta \ln c_{t+1} + [(\varepsilon-\sigma)/\sigma(\varepsilon-1)]\Delta \ln \alpha_{t+1}}$ and linearize $\Delta \ln \alpha_{t+1}$ around the point $z_{t+1} = z_t$, which delivers

$$M_{t+1} \approx \beta \exp\left(-\frac{1}{\sigma} \Delta \ln c_{t+1} + \frac{(\varepsilon - \sigma)}{\sigma(\varepsilon - 1)} (1-\alpha_t) \Delta \ln z_{t+1} \right). \tag{10}$$

This formula summarizes how composition risk affects pricing. First, agents want to be compensated for correlation of returns with the "composition risk factor" $\Delta \ln z_{t+1}$, the change in the consumption-housing expenditure ratio. Second, the importance of

[7] This is a symmetric property: equivalently housing services are more valued in states where the expenditure share of housing increases.

composition risk decreases as the share of non-housing consumption α_t increases. As $\alpha_t \to 1$, the pricing kernel becomes more and more similar to the standard CCAPM. Third, the coefficient $(\varepsilon - \sigma)/\sigma(\varepsilon - 1)$ determines what the 'bad states' are. With a positive composition effect, agents fear states in which the consumption-housing ratio increases, and they prefer assets which pay out in those states. The converse is true with a negative composition effect. The composition effect thus determines whether composition risk increases or decreases premia relative to the CCAPM.[8]

To derive the behavior of returns, we now impose some structure on the vector $(\Delta \ln c_t, \ln z_t)$. We assume that the vector follows a stationary VAR with conditionally normal errors. In other words, the log expenditures on consumption and housing are cointegrated. We impose the following restrictions on this VAR. Consumption growth is i.i.d. just like in Campbell and Cochrane (1999). In particular, expected consumption growth is constant:

$$\Delta \ln c_{t+1} = \mu_c + u^c_{t+1}, \tag{11}$$

where the consumption growth shock u^c_{t+1} has mean zero and variance v_c.

The log expenditure ratio is an autoregressive process

$$\ln z_{t+1} = (1-\rho)\mu_z + \rho \ln z_t + u^z_{t+1}, \tag{12}$$

where u^z_{t+1} has mean zero and conditional variance $v_{z,t}$. The shocks u^c_{t+1} and u^z_{t+1} are conditionally normal and uncorrelated. The key properties of this process are that (i) consumption growth is not forecastable, (ii) the consumption housing ratio is persistent, but past consumption growth does not help forecast it, (iii) there is heteroskedasticity in disturbances to the log expenditure ratio, and (iv) disturbances are conditionally normal. We will provide evidence for all these features in the next section. Here we focus on the qualitative implications.

3.5 Behavior of the Riskless Rate

Using conditional normality of $(\Delta \ln c_{t+1}, \ln z_t)$, we can obtain a simple approximate formula for the log interest rate based on the pricing kernel equation (10):

$$\begin{aligned} r^f_{t+1} \approx & -\log\beta + \frac{1}{\sigma}\mu_c - \frac{1}{2\sigma^2}v_c \\ & + (1-\alpha_t)\left\{ -\frac{(\varepsilon-\sigma)}{\sigma(\varepsilon-1)}(\rho-1)(\ln z_t - \mu_z) - \frac{1}{2}\left[\frac{\varepsilon-\sigma}{\sigma(\varepsilon-1)}\right]^2 (1-\alpha_t)\, v_{z,t} \right\} \end{aligned}$$

[8] The approximation is not exact. In particular, it masks the fact that in the nonlinear kernel, correlation with consumption growth will be also be weighted differently as α_t changes. This effect could be made explicit by picking a different linearization point. For example, one could assume an AR(1) process for $\ln z_t$ and linearize around the conditional mean. We choose not to do this here, because the current approximation is simpler and is sufficient to interpret the computational results, which are based on the true nonlinear kernel.

The effect of beliefs about consumption growth is familiar. If consumption is expected to grow, agents try to borrow, pushing the interest rate up. If consumption growth becomes more uncertain, agents try to engage in precautionary saving, pushing the interest rate down. Both effects are stronger the more agents want to smooth consumption (low σ). Since consumption is not very volatile, the precautionary savings effect is small for reasonable values of σ; this is the riskfree rate puzzle.

The term in braces represents the impact of composition risk on the interest rate. First, a positive composition effect decreases the mean interest rate. *Precautionary saving motivated by composition risk thus helps resolve the riskfree rate puzzle.* Second, time variation in interest rates is governed by both consumption-smoothing and precautionary-savings effects. The former effect arises because when $\ln z_t$ is high, agents expect the consumption-housing share to go down ('good states to come'). This leads agents to borrow more and push up the interest rate. It thus makes the interest rate increase with the distance of $\ln z_t$ from its unconditional mean.

The precautionary savings effect is more subtle. On the one hand, the impact of composition risk is also diminished as the non-housing share rises. Indeed, as $\alpha_t \to 1$, the precautionary savings effect will vanish faster than the consumption-smoothing effect. This implies that the interest rate for high α_t will be higher than the mean. At least for high α_t, we can thus expect an interest rate function that is increasing in α_t. On the other hand, there may be a counteracting effect if the conditional variance $v_{z,t}$ increases in $\ln z_t$. We will provide evidence for such an effect in Section 4. This increase in risk tends to make the interest rate lower as $\ln z_t$ increases. Heteroskedasticity can thus help explain why the riskless rate need not be highly volatile, even though the stochastic discount factor depends on the persistent state variable α_t.

3.6 Time Varying Risk Premia

Using conditional normality of $(\Delta \ln c_{t+1}, \ln z_t)$, we obtain an approximate formula for the expected return r on an asset in excess of the riskless rate:

$$\begin{aligned} E_t\left(r_{t+1}\right) - r^f_{t+1} + \frac{1}{2} var_t(r_{t+1}) \approx\ & \frac{1}{\sigma} cov_t\left(\Delta \ln c_{t+1}, r_{t+1}\right) \\ & - \left(1-\alpha_t\right) \frac{(\varepsilon - \sigma)}{\sigma\left(\varepsilon - 1\right)} cov_t\left(\Delta \ln z_{t+1}, r_{t+1}\right). \end{aligned} \tag{13}$$

This equation illustrates how conditional risk premia depend on the two sources of risk, consumption and composition risk. To provide intuition, we interpret it in terms of market prices of risk. Of course, this interpretation has to be taken in with some caution, since the approximation masks nonlinearities in the pricing kernel. Following standard language, we call the covariance of returns with consumption growth the quantity of consumption risk, while the price of risk is the coefficient of relative risk aversion, $\gamma = 1/\sigma$.

For composition risk, it is intuitive to define the quantity of risk as the correlation of *relative* quantities with returns. For the relative quantity of nondurables to shelter, this is approximately equal to $\frac{\varepsilon}{\varepsilon-1}cov_t\left(\Delta \ln z_{t+1}, r_{t+1}\right)$. The volatility of relative quantities is high, for given volatility of relative expenditure, when the intratemporal elasticity of substitution is close to one, that is, when preferences are close to Cobb-Douglas. The 'price of composition risk' can then be defined as $1/\sigma - 1/\varepsilon$. It is increasing in the risk aversion parameter. More generally, it is large when agents prefer to substitute across goods rather than over time. Another property is that the whole composition effect becomes less important the lower the share of housing (higher α_t).

It is well known that the conditional covariance of returns and consumption growth is small and does not vary much over time. This implies that the CCAPM generates a low equity premium and little predictability. Composition risk may be helpful on both counts. It will matter if either its quantity is large and variable or if its price is high. Its quantity is high and variable if the covariance of returns and $\Delta \ln z_t$ varies a lot *or* if preferences are close to Cobb-Douglas. Its price is high if risk aversion is high. In what follows, we provide a quantitative assessment.

4 Data

In this section, we introduce the data used in our empirical work. We discuss various measurement issues and present a VAR for consumption growth and the consumption-housing expenditure ratio that supplies the forcing process for our equilibrium calculation.

4.1 Consumption Data

The two important variables in our model are non-housing consumption c and its expenditure share α. We measure these variables using aggregate data from the NIPA tables. These tables contain consumption and expenditure data subdivided into different categories. NIPA treats housing differently from other durable goods, in that the tables contain data on the service flow of housing and expenditures on the service flow. For other durable goods, such as cars, the tables only contain data on the good itself, not its service flow.

To measure non-housing consumption c, we follow the convention and use aggregate consumption of nondurables and services. However, we deviate from the convention when we *exclude housing services*. Table 1 shows that our measure of c grows at an average rate of 2.2% and varies only little, 1.9% per year. For comparison, we report the corresponding numbers for the conventional consumption growth measure which includes housing in the last column of Table 1. When we include housing services, consumption growth is a little higher on average and somewhat less volatile. Non-housing consumption growth is positively autocorrelated, especially during the postwar sample. Again, this feature is also present when we include housing consumption.[9]

[9] Consumption growth is negatively autocorrelated over the sample period 1889 to 1978. Mehra and

Table 1 also shows that consumers spend on average 82.6% of their total expenditure on non-housing consumption. This expenditure share varies little over time, only 1.5%. In other words, consumers spend about the same fraction of their total expenditures on non-housing consumption over time, apart from some large movements. Figure 1 shows that these large movements happened, for example, during the Great Depression and right after WWII. The non-housing expenditure share α takes on its lowest value of 77% in 1932, and its highest value of 87% in 1946. Together with a high autocorrelation of 0.965, low variations in α translate into the low frequency movements that we see in Figure 1. This aggregate time series evidence from the NIPA tables is consistent with cross-sectional microdata. The publication "Housing Expenditures" by the U.S. Department of Labor documents the housing expenditure share for different types of households using the Consumer Expenditure Survey conducted in 1999. Households that differ in the number of people, age, urban or rural living location, and other characteristics still spend the same percentage, 19%, of their expenditures on shelter.

Consumption data are affected by the usual measurement problems. These problems are particularly severe when it comes to housing. In constructing the NIPA tables, the Bureau of Economic Analysis combines data on rents from various sources, including the CPI, and estimates of the residential housing stock. We work mostly with expenditure data, which we view as more reliable than prices or quantities of housing consumption. The main advantage of expenditure data is that α does not seem to trend over time, as we see from Figure 1. Figure 2 plots the log of real rents q and relative quantities s/c. The two series trend in opposite directions. Housing services have become cheaper over time relative to other nondurables and services. As the within-period first order condition (7) would predict, more housing services were consumed. Since 1982, these trends are reversed.

Prescott (1997) argues that the recent upward trend in relative rents is due to measurement errors associated with housing services and has nothing to do with the costs of housing services, which is what we want to measure. These measurement errors have several sources. For example, the market prices for most goods, such as milk, can be obtained using store surveys. There are, however, no market prices for owner-occupied housing. The Bureau of Labor Statistics therefore asks home owners for an estimate of how much their house would rent for. These estimates introduce errors in rent data. But even if rents were measured correctly for these home owners, they would not reflect the cost of housing services. The true costs of housing services depend on many other factors that are ignored in the construction of the NIPA data. Such factors include who rents the house, who owns the house, the tax situation and financing method of the owner, and the moral hazard problems associated with renting.

We use annual data, because it is available going back to 1929, instead of the short postwar quarterly sample. Figure 1 shows that the non-housing expenditure share and the dividend-yield moved sharply in opposite directions, while the two series are positively correlated overall. This episode thus must be driven by shocks that are outside of our

Prescott (1985) compute an autocorrelation coefficient of -0.14 over this period. Over this long sample, consumption growth volatility is also higher, 3.6% per year.

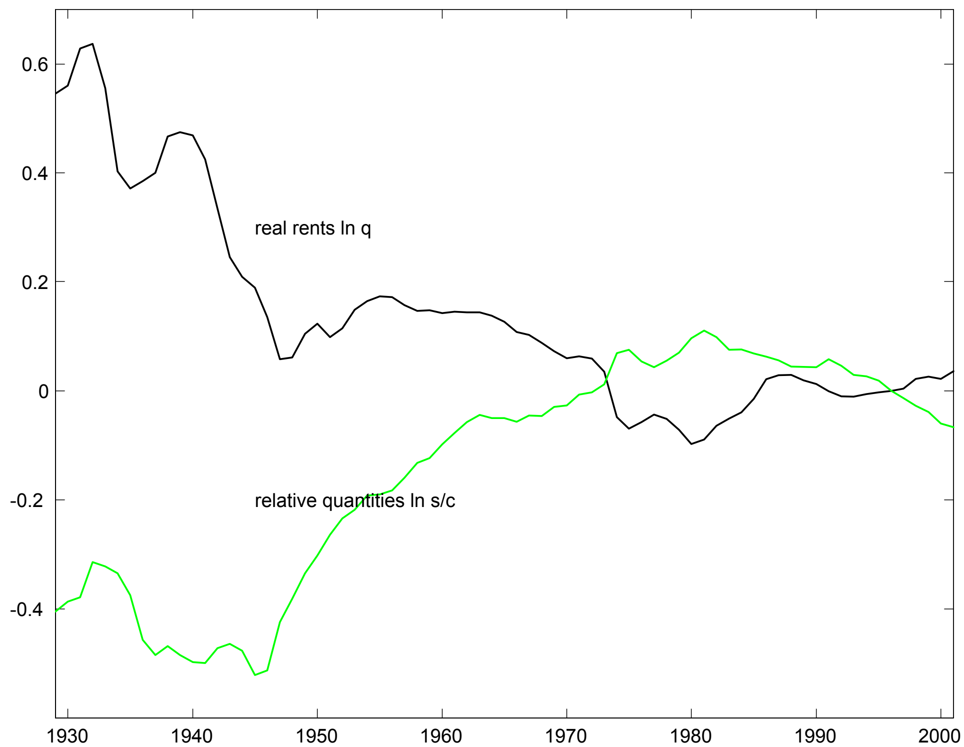

Figure 2: Annual NIPA data on real rents q and the quantity s of housing services relative to the quantity c of nondurables and non-housing services.

model. We therefore start our sample only in 1936. In doing so, we omit a period of highly volatile consumption growth, which would help our consumption-based explanation of asset prices.

4.2 Asset data

Real housing returns are computed using NIPA tables. We take the real housing value $p_t^h h_t$ from the NIPA Fixed Asset Tables 2.1, line 68. This series computes the nominal housing value using the "current value method" which measures the year-end market value of residential housing structures. To include the value of land, we assume that land prices are perfectly corrected with the price of structures. Using Census data, we estimate that the value of the land is 36% of the total housing value. We therefore adjust houses prices to $p^h/(1-0.36)$. The dividends on housing are rent payments during that year, $q_t s_t$. From these dividends, we subtract real depreciation $\delta p_{t-1}^h h_{t-1}$, where we fix δ at 1.6% following Davis and Heathcote (2001). We follow Flavin and Yamashita (2002) and subtract a net real property tax payment of $(1-0.33) \times 0.025 \times p_{t-1}^h h_{t-1}$, where the marginal tax rate is assumed to be 33% and the property tax rate is assumed to be 2.5%.

The real housing return is thus

$$\frac{\left(p_t^h h_t + q_t s_t\right)/h_t}{\left(p_{t-1}^h h_{t-1}\right)/h_{t-1}} - \delta - (1 - 0.33) \times 0.025. \tag{14}$$

In the appendix, we compare our housing return series with alternative measures of aggregate housing returns.

Nominal stock prices and corresponding dividends, and the nominal riskfree rate are from Robert Shiller's website. To deflate these returns, we construct the price index that corresponds to our measure of consumption from the NIPA numbers.

Table 1 shows that housing pays average returns somewhere between those on stocks and the riskfree rate. The volatility of housing is phenomenally low. While the volatility of stock returns is about 15-17%, housing returns are about as volatile as the riskfree rate, 2-3%! These numbers do not even differ much over subsamples, except that the volatility of housing and the real rate was even 1% lower in the postwar period. Finally, the autocorrelation of real returns on housing is 48%, huge compared to almost zero for stocks. The high autocorrelation makes housing look even more attractive compared to stocks.

Do these numbers mean that houses an incredible bargain? No. Returns on individual houses are much more volatile than returns in the aggregate. For example, Flavin and Yamashita (2002) report standard deviations of real returns for individual houses of 14% computed using PSID data from 1968 to 1992 in their Table 1A. Housing returns based on aggregate housing indices are much less volatile, even at disaggregated regional levels. For example, Table 1B in Flavin and Yamashita reports that real housing returns based on the Case and Shiller (1989) data on house prices in 4 metropolitan areas are only 1/3 to 1/2 as volatile as on the individual level. The appendix reports similar ratios for state-level housing returns. More evidence is discussed in Caplin et al. (1997). Thus, it seems that most of the volatility in housing returns is due to variations in the price of individual houses, as opposed to price movements on the level of cities, regions or nations.

Table 1: Summary statistics

	$\Delta \ln c$	α	$\Delta \ln z$	r^s	r^h	r^f	conventional $\Delta \ln c$
mean (%)	2.17	82.6	0.05	6.94	2.52	0.75	2.25
auto	0.23	0.965	0.60	-0.06	0.48	0.73	0.24
				Postwar sample			
mean (%)	1.85	82.3	-0.63	7.80	2.09	1.57	1.98
auto	0.40	0.84	0.67	0.02	0.44	0.52	0.41
				Covariances ($\times 10^{-4}$)			
$\Delta \ln c$	$(1.88)^2$						$(1.58)^2$
α	0.08	$(1.54)^2$					0.18
$\Delta \ln z$	0.50	0.10	$(0.40)^2$				0.40
r^s	1.30	-0.58	-1.43	$(16.56)^2$			0.08
r^h	2.73	0.42	0.10	0.43	$(2.73)^2$		2.27
r^f	0.11	-4.00	-0.78	12.06	0.17	$(3.68)^2$	-0.32
				Postwar sample (Covariances ($\times 10^{-4}$)			
$\Delta \ln c$	$(1.46)^2$						$(1.23)^2$
α	-0.85	$(1.28)^2$					-0.63
$\Delta \ln z$	0.18	-0.20	$(0.36)^2$				0.12
r^s	3.26	0.29	-1.35	$(15.36)^2$			1.79
r^h	1.64	-0.16	0.09	1.84	$(2.34)^2$		1.39
r^f	1.74	-2.30	0.03	6.16	-1.09	$(2.86)^2$	1.01

NOTE: Summary statistics are computed over the whole sample 1936 to 2000, and over the postwar sample 1947-2000. Non-housing consumption data $\Delta \ln c$ is based on NIPA Table 7.4. The data is nondurables (line 6) and services (line 13), minus shoes and clothing (line 8) and housing services (line 14). The non-housing consumption expenditure share α is based on NIPA Table 2.2. The data on non-housing expenditures is from the same lines as non-housing consumption, and housing expenditures are from line 14. The expenditure ratio z is $\alpha/(1-\alpha)$. Log real stock returns r^s and the log real rate r^f are from Robert Shiller's website. Log real housing returns r^h are constructed as in equation (14) from NIPA data. We construct a price index using data from Tables 2.2 and 7.4 corresponding to our definition of non-housing consumption c to deflate returns. The last column indicates the conventional measure of consumption growth which includes housing (line 13, Tables 2.2 and 7.4). Covariance numbers with real returns in the last column are deflated with the deflator corresponding to the conventional consumption measure.

5 Evidence about First-order Conditions

5.1 Static First-order Condition

Taking logs of the static first-order condition, we can derive a cointegrating relationship, which we use to obtain estimates of the intratemporal elasticity of substitution ε. Table 2 reports the results of the Johansen-test for cointegration of log real rents $\ln q$ and the log quantity ratio $\ln s/c$. The test allows for linear trends in the data, and includes 2 lags. We strongly reject the null of no cointegration. The estimate of ε implied by the estimated cointegrating equation is greater than 1, indicating that housing services s and nondurable consumption c are substitutes. The standard errors indicate that ε is not likely to be below one. In other words, we find that the utility function is not likely to be Cobb-Douglas. These results do not change much when we extend the sample back to include the Great Depression. The immediate postwar period leads to estimates of ε which are close to but below 1. The standard errors computed over this period are, however, large (not reported).

Table 2: Estimate of intratemporal elasticity

LR		ε	
21.75	[20.04]	1.27	(0.16)

NOTE: The first row reports the likelihood ratio of the Johansen-test for cointegration and the corresponding 1% critical value in square brackets. The second row reports ε from the cointegrating equation $\ln s_t/c_t + \varepsilon \ln q_t + \text{constant}$ and the standard errors in round brackets. Sample period: 1936-2001.

These results confirm those in Ogaki and Reinhart (1998) who estimate the elasticity of intratemporal substitution between nondurables and durables (not housing). Their Table 2 on page 1091 gives $[1.04, 1.43]$ as a 95% confidence interval for ε. The test for unitary elasticity $\varepsilon = 1$ is thus rejected. Mayo (1981) estimates a value for ε between 0.3 and 0.9. Hanushek and Quigley (1980), however, argue that housing demand may not respond much to a given price change in the presence of adjustment costs. Estimates of price elasticities measured over short horizons may therefore be biased downwards substantially. Hanushek and Quigley provide evidence based on regional house price indices that this bias may be substantial.

We do not estimate the elasticity of intertemporal substitution σ. Hall (1988) estimates σ to be around 0.2. Ogaki and Reinhart (1998) report estimates between $[0.32, 0.45]$. Studies based on micro data typically find somewhat higher values for σ. Runkle (1991) reports an estimate of 0.45 using micro data on food consumption. Attanasio and Browning (1995) report estimates using Consumer Expenditure Survey data between [0.48,0.67].

To sum up, these empirical studies are consistent with (i) ε close but different from 1, and (ii) σ much smaller than ε, and not larger than 0.68. In other words, the utility function is close but not equal to Cobb-Douglas. Moreover, the intratemporal elasticity of substitution between nondurables and durables ε is much larger than the intertemporal elasticity of substitution σ. From this evidence, we conclude that

$$\sigma < 1 < \varepsilon.$$

5.2 Euler equation

In this section, we begin our exploration of the role of composition risk by comparing unconditional mean returns and unconditional covariances of returns with our stochastic discount factor. Any (gross) return R_{t+1} must satisfy

$$E\left[R_{t+1}\right] E\left[\frac{1}{R^f_{t+1}}\right] - 1 = -cov\left(M_{t+1}, R_{t+1}\right),$$

where R^f is the (gross) riskfree rate. The left hand side of this equation is approximately equal to the average excess return. Table 3 shows the right hand side of this equation for stock and house returns. The model implied premia are computed as a function of the intratemporal elasticity ε, holding the coefficient of relative risk aversion fixed at $\gamma = 1/\sigma = 5$ and the discount factor at $\beta = 0.99$. The value of ε is chosen to be larger than σ, consistent with the evidence discussed in Section 5. The numbers are not sensitive to the value of β.

Table 3: Euler-equation Implications (in %)

ε	1.1	1.05	1.01
equity premium	0.75	1.35	8.58
housing premium	0.11	0.13	6.94

NOTE: The table entries are $-cov\left(M_{t+1}, R_{t+1}\right)$ computed using the pricing kernel (2) for $\beta = 0.99, \gamma = 1/\sigma = 5$ and the values of ε indicated in the top row. The sample period is 1936-2000.

To illustrate the contributions of consumption and composition risk, it is convenient to linearize the discount factor M_{t+1} around the point $(\Delta \ln c_{t+1}, \ln \alpha_{t+1}) = (0, \ln \alpha_t)$. To first order, we obtain.

$$M_{t+1} \approx \beta\left(1 - \frac{1}{\sigma}\Delta \ln c_{t+1} + \frac{(\varepsilon - \sigma)}{\sigma\left(\varepsilon - 1\right)}\Delta \ln \alpha_{t+1}\right) \tag{15}$$

According to this approximation, our model behaves like a linear two-factor model, with consumption growth and the change in the log non-housing share as factors. Excess returns can thus be decomposed into premia due to the two factors:

$$E[R_{t+1}]E\left[\frac{1}{R_{t+1}^f}\right] = \beta\frac{1}{\sigma}cov(\Delta\ln c_{t+1}, R_{t+1}) - \beta\frac{(\varepsilon-\sigma)}{\sigma(\varepsilon-1)}cov\left(\Delta\ln\alpha_{t+1}, R_{t+1}\right).$$

The standard model without housing needs to entirely rely on the first term in this Euler equation to generate excess returns. Tiny values of the intertemporal elasticity σ are needed in this model to generate large equity premia given the low covariance of returns with consumption growth. In the model with housing, a positive composition effect increases the equity premium beyond the one implied by the standard model. From Table 1, we see that the covariance of stock returns with expenditure growth $\Delta\ln\alpha_{t+1}$ is negative. Now, if intertemporal smoothing is important ($\varepsilon > \sigma$) and consumption and housing are substitutes ($\varepsilon > 1$), the second term in the Euler equation is positive. As the utility function approaches Cobb-Douglas (as ε approaches 1), the equity premium becomes large.

We go on to illustrate how the conditional covariance between returns and the pricing kernel changes over time. Since a conditional version of the Euler equation holds,

$$E_t\left[R_{t+1}\right]\frac{1}{R_{t+1}^f} - 1 = -cov_t\left(M_{t+1}, R_{t+1}\right),$$

the time variation in the covariance captures the predictability in excess returns. Figure 3 shows a rolling window estimate of the conditional covariance $-cov_t\left(M_{t+1}, R_{t+1}\right)$. The window we use is 5 years and is centered around the current period. We can see that the covariance of the pricing kernel M and returns varies over time and is roughly able to capture expected excess stock returns for negative ϕ. For example, the conditional covariance is high in 1960s, early 1970s, mid 1980 when stocks did well, and low in the early 1950s, early 1980s when stocks did poorly.

6 Equilibrium Prices

In this section, we consider asset pricing in a model economy in where housing and numeraire consumption shocks are the only sources of uncertainty. We show that there exists a model economy which does a reasonable job in matching both asset price dynamics and the dynamics of consumption growth and house shares.

6.1 Matching Model and Data

To price equity, we follow Campbell (1986) and Abel (1994) and assume that dividends d_t equal a 'levered up' version of consumption

$$\Delta\ln d_t = k\Delta\ln c_t$$

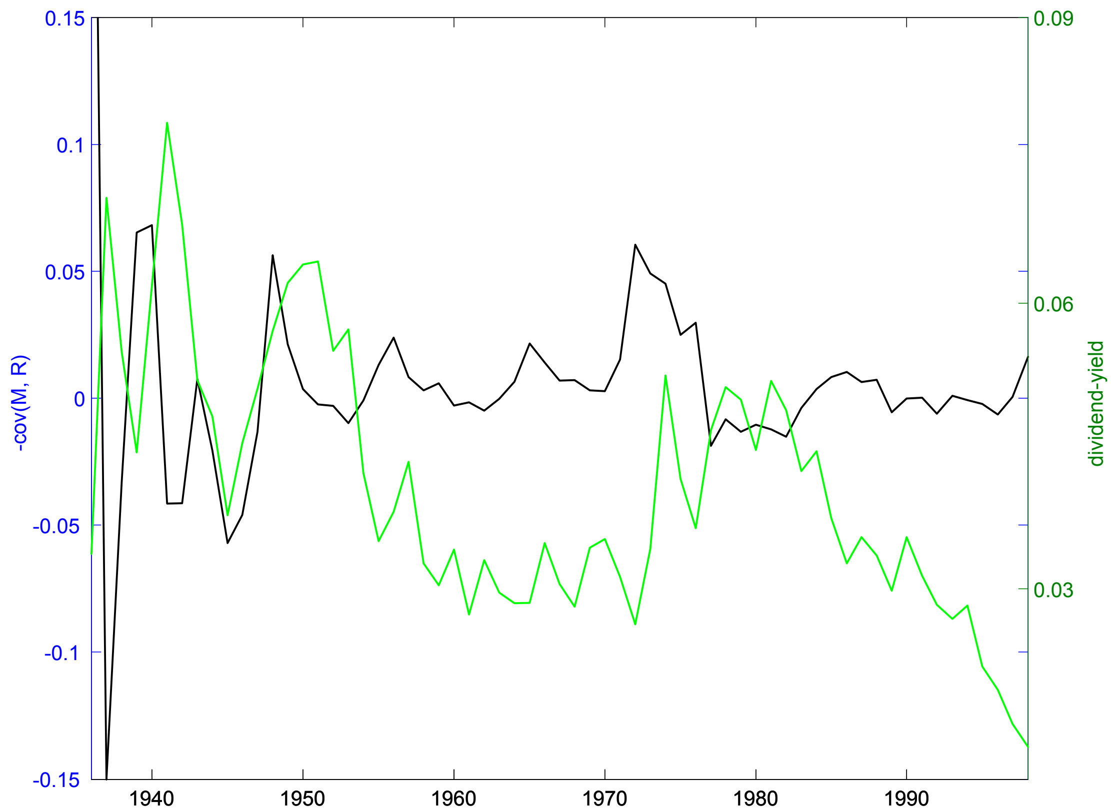

Figure 3: Dividend yield and rolling window estimate of $-cov_t(M_{t+1}, R_{t+1})$ for real stock returns, computed using a window is 5 years which is centered around t. The pricing kernel M is computed using $\varepsilon = 1.05$ and $\gamma = 1/\sigma = 5$.

Dividends equal consumption for $k = 1$. With $k > 1$, this specification captures levered equity. We follow Abel (1994) in using $k = 2.74$. This value is calibrated to the volatility of consumption growth compared to dividend growth. A difference equation for the price-dividend ratio is obtained by plugging the discount factor (1) into the pricing equation. For housing, we calculate an analogous price-dividend ratio by equating the value of the housing stock with the present discounted value of all future housing services, $q_t s_t = c_t \exp(-\ln z_t)$:

$$\begin{aligned} \upsilon_t^s &= \frac{p_t^s}{d_t} = E_t\left[M_{t+1}\left(\upsilon_{t+1}^s + 1\right)e^{k\Delta \ln c_{t+1}}\right] \\ \upsilon_t^h &= E_t\left[M_{t+1}\left(\upsilon_{t+1}^h + 1\right)e^{\Delta \ln c_{t+1} - \Delta \ln z_{t+1}}\right] \end{aligned} \tag{16}$$

Conveniently, the solution υ_t^s reduces to the price of a consol bond if $k = 0$.

According to our definition of equilibrium, prices are functions of the 'fundamentals' process $\left(y_t, \overline{h}_t\right)$. For a given initial condition y_0, we can also express prices as functions of the process $(\Delta \ln c_t, \ln z_t)$. We estimate a bivariate heteroskedastic VAR for $(\Delta \ln c_t, \ln z_t)$ as the forcing process. Given a stationary process $(\Delta \ln c_t, \ln z_t)$, the price-dividend ratios for equity and housing as well as the consol price can be determined and are also stationary. In addition, if α_t is observed, all asset prices and returns of interest can be determined without knowing the parameter ω. Indeed, while the level of housing dividends depends on ω, its growth rate does not. In what follows, we thus consider model economies that are determined by a $(\Delta \ln c_t, \ln z_t)$-process as well as elasticities σ and ε, discount factors β and leverage parameters k and compare their properties to the data.

Parameters for the consumption growth process can be taken from Table 1. The consumption growth process is autocorrelated. Heaton (1993, 1995) argues that this autocorrelation may be entirely due to time aggregation. We therefore assume that expected consumption growth μ_c is constant and set it to 2.17%. We also specify the variance of consumption growth v_c to be constant, and set it equal to 0.0188^2.

A regression of $\ln z_{t+1}$ on its lagged value and $\Delta \ln c_t$ shows that consumption growth is barely significant. We therefore specify the log expenditure share as the autoregressive process. The shocks u_t^z to the log expenditure ratio show substantial heteroskedasticity. In particular, their variance seem to increase with $\ln z_t$. To capture this heteroskedasticity, we specify the conditional variance as

$$v_{z,t} = a_1^z \max\{\ln z_t, \overline{z}\} - a_0^z.$$

Under this specification, the conditional variance is linear in $\ln z_t$ except for small $\ln z_t$, where it is constant. Table 4 reports the parameter estimates and t-statistics. The estimation is in two steps. The first step estimates μ_z and ρ in equation (12) using ordinary least squares and saves the squared residuals. The second step regresses squared residuals on a constant and $\ln z_{t-1}$ to estimate a_0^z and a_1^z. The precise value for $\overline{z}$ does not matter much in our application. We fix the value to match the unconditional variance of $\ln z$.

Table 4: Estimates of expenditure ratio dynamics

μ_z	ρ	a_0^z	a_1^z	$\overline{z}$
1.56	0.96	−0.0117	0.0081	1.47
(52.29)	(13.20)	(−3.82)	(3.92)	

NOTE: The parameters are estimated in two steps, as explained in the text. T-statistics in brackets are based on 4 Newey-West lags. T-statistics of the second step do not take into account estimation risk from the first step. Sample period: 1936-2000.

Figure 4 shows the empirical and simulated density of the log expenditure share $\ln z_t$. The empirical density is skewed to the left, and this skewness is captured by the estimated data-generating process in Table 4, as the simulated pdf shows.

For given VAR parameters μ_c, μ_z and ρ, conditional variance parameters v_c, s_0^z, and s_1^z, and a given a choice of the preference parameters ω and ε, the actual exogenous forcing process $(y_t, \bar{h}_t)$ could be backed out using (8) and an initial condition y_0.

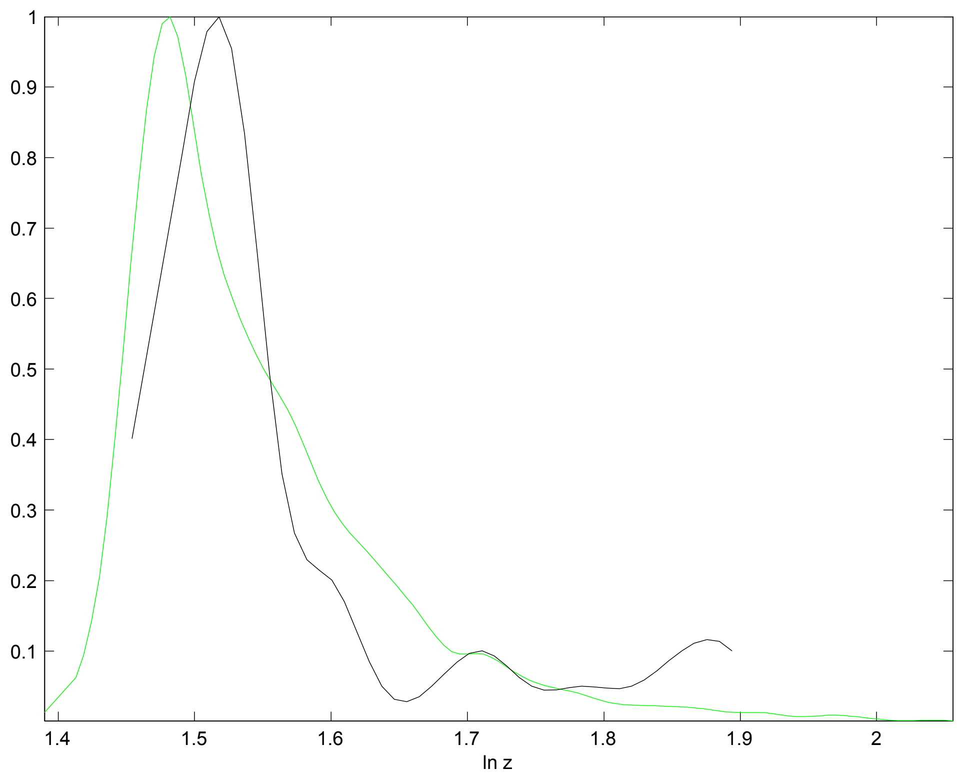

Figure 4: Empirical density of the log expenditure share $\ln z$ (black line) and simulated density (gray line).

6.2 Results

Table 5 reports first and second moments of annual equity premia, consol premia, riskless rates and housing returns, all in logs, for a set of parameters. The results are based on averages over 5000 simulations of length 65, the length of our data sample, all starting at the consumption housing ratio for 1936. If there is enough composition risk (that is, ε is close enough to one), the model matches, with a relatively low risk aversion coefficient of 5, a high and volatile equity premium together with a small and smooth consol premium and riskless rate. In addition, the model does a reasonable job on housing returns.[10] The bold faced case, $\varepsilon = 1.05$, is our leading example and will now be described in more detail.

[10] Housing returns computed from the model according to equation (16) not only measure the returns on one unit of housing, but also include the value of new housing. Housing returns computed for Table 1 do not include the value of new housing. When we do include the value of new housing, mean returns go up by 2.5 percentage points. The standard deviation of housing returns is unchanged.

Table 5: Model Implications (%) for $\beta = .99$.

γ	ε	$E\left[r^f\right]$	$E\left(er^s\right)$	$E\left(er^b\right)$	$E\left(er^h\right)$	$\sigma\left(r^f\right)$	$\sigma\left(er^s\right)$	$\sigma\left(er^b\right)$	$\sigma\left(er^h\right)$
5	1.25	11.1	1.1	0.3	0.8	0.7	5.9	2.0	5.0
	1.10	9.0	1.8	0.8	1.5	0.9	7.6	3.4	6.8
	1.05	**1.8**	**5.8**	**2.5**	**3.7**	**1.8**	**14.7**	**5.9**	**10.1**
	1.04	−2.3	9.22	4.4	5.6	1.8	20.9	9.1	13.6
3	1.05	5.2	2.7	1.1	1.6	0.9	12.6	5.3	9.0
4		4.2	3.9	1.7	2.7	0.6	13.2	5.5	9.7
6		−1.9	8.8	4.1	5.6	2.2	18.3	4.1	12.5
7		-4.1	10.6	5.1	6.2	3.7	19.7	10.1	13.8

Figure 5 plots asset prices as a function of the single state variable, the consumption-housing ratio $\ln z_t$. All prices of long-lived assets are decreasing in $\ln z_t$, with stock prices showing the most sensitivity.[11] Since consumption growth is not forecastable, changes in prices cannot be due to changes in expected dividend growth. Instead, all price movements must be due to either changes in discount rates or changes in dividend volatility. Using the approximation (10), the state-dependent discount rate is

$$\delta_{t+1} = -\log M_{t+1} = -\log \beta + \gamma \Delta \ln c_{t+1} + (1-\alpha_t) \frac{(\varepsilon - \sigma)}{\sigma(\varepsilon - 1)} \Delta \ln z_{t+1}$$

Price Volatility: the Role of News about Discount Rates

We first illustrate the role of discount rate changes in the absence of heteroskedasticity. We use the approximate pricing kernel (10) to rewrite the price-dividend ratio (16) as

$$\begin{aligned} \upsilon^s(\ln z_t) &= E_t\left[e^{-(\delta_t - k\Delta c_{t+1})}\left(1+\upsilon^s(\ln z_{t+1})\right)\right] \\ &= E_t\left[\sum_{j=1}^{\infty} \exp\left(-\sum_{i=1}^{j}\delta_{t+i}\right) e^{k(c_{t+j}-c_t)}\right] \\ &= \sum_{j=1}^{\infty} \beta^j e^{\left((k-\gamma)\mu_c + \frac{1}{2}(k-\gamma)^2 v_c\right)j} E\left[\exp\left(-\frac{(\varepsilon-\sigma)}{\sigma(\varepsilon-1)}\sum_{i=1}^{j}(1-\alpha_{t+i})\Delta \ln z_{t+i}\right)\right] \\ &= \sum_{j=1}^{\infty} \upsilon_j^*(k)\, w_j(\ln z_t). \end{aligned}$$

In this formula, $\upsilon_j^*(k)$ is the present discounted value, adjusted for consumption risk and normalized by current dividends, of a claim to dividends in period $t+j$. In other words, it is the value of such a claim in the standard Lucas model without housing. With i.i.d. consumption, this benchmark price is constant, which amounts to a version of the "volatility puzzle", emphasized by Campbell (2000). In our model, volatility is induced

[11] Kinks in the price function for low values of $\ln z_t$ occur since the conditional variances of the innovations become constant as $\ln z_t$ drops below $\bar{z}$. The results are not sensitive to the choice of $\bar{z}$.

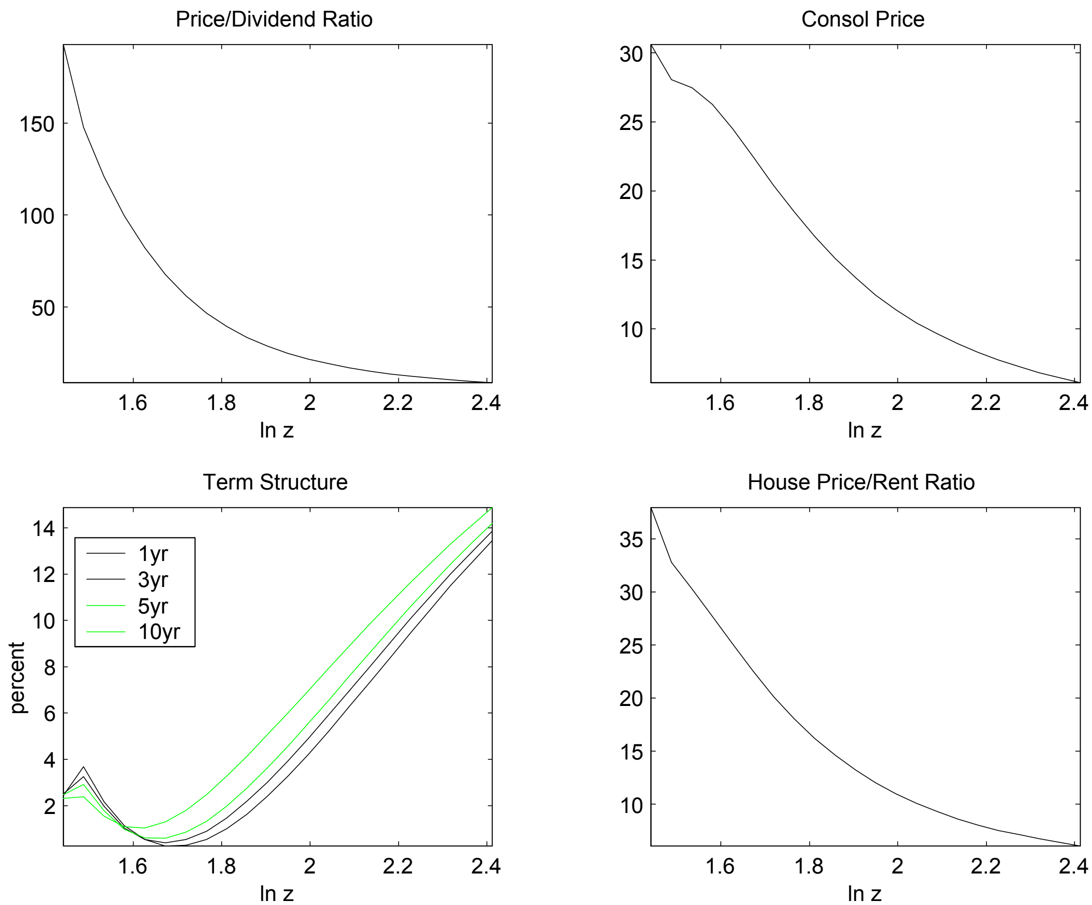

Figure 5: Asset prices as a funciton of $\ln z$.

by the new *discount factor for composition risk*, $w_j(\ln z_t)$. Innovations to $\ln z_t$ provide news about current and future discount rates δ_{t+j}. In line with the empirical findings of Cochrane (1994), it is this type of news, not news about dividends, that accounts for most changes in prices.

The above formula clarifies some properties of stock and consol prices. First, as long as v_j^* is increasing in k, the price-dividend ratio is larger and more volatile than the consol price. The random factor $w(\ln z_t)$ affects consols and stocks in the same way; differences across assets exist only to the extent that there are differences in $v_j^*(k)$. The latter is increasing if $\mu_c > (\gamma - k)\, v_c$. An increase in k increases both mean growth and dividend risk. If risk aversion and consumption risk are not too high, the mean effect dominates and the price-dividend ratio goes up. Second, mean reversion in the consumption-housing ratio z_t implies that both prices move opposite to z_t. Unfortunately, the discount factor $w(\ln z_t)$ is not available in closed form, since z_t itself is a nonlinear function of α_t. However, numerical results for the homoskedastic case (not shown) deliver a shape much like Figure (5). If z_t is large, then $\Delta \ln z_{t+1}$ is negative with high probability, that is, housing is expected to grow faster than expenditure on numeraire. This makes saving in numeraire terms relatively less attractive and hence lowers asset prices.

Price Volatility: The Role of Heteroskedasticity

With heteroskedasticity, an increase $\ln z_t$ also increases risk, as $v_{z,t}$ increases with $\ln z_t$. This has two effects. First, it dampens the response of prices to a unit change in the consumption-housing share $\ln z_t$. Indeed, an increase in risk encourages precautionary savings, that is, agents discount the future less, so that prices fall by less than in the homoskedastic case. At the same time, the size of the typical shock is larger when $\ln z_t$ is high. This tends to increase the conditional volatility of returns when $\ln z_t$ is high. On net, the variance of returns and their associated premia tend to be higher in the heteroskedastic case.

Another important effect of heteroskedasticity is to reduce the volatility of the riskless rate. Comparing Figure 4 and Figure (5), it is apparent that the riskless rate is very stable in the part of state space which has highest probability. Indeed, it is nonmonotonic in this area, as a result of the counteracting precautionary savings and consumption smoothing effects.

Time-Varying Risk Premia

Since prices move opposite to the composition risk factor $\ln z_t$, returns will be negative correlated with $\Delta \ln z_t$. According to the risk premium equation (13), this is exactly what is needed to generate an equity premium higher than that plausibly implied by consumption risk. Since the effect of discount rates on stocks and consols is qualitatively similar, our model also predicts a positive premium on consols. However, since the stock holders are much more exposed to changes in discount rates, due to both growth in dividends and leverage, the equity premium is much larger and more volatile than the consol premium.

The model also implies that excess returns are predictable. This is illustrated in Figure 6, which plots the model-implied dividend yield and expected returns over the post-war period. It is apparent that the dividend yield is a slow-moving state variable which forecasts returns. This is consistent with recent empirical evidence. Moreover, the model predicts that α_t, which is highly correlated with the dividend yield should also be a good forecasting variable. This is indeed the case in the data, as documented in the next section.

Housing Stock Value

Figure 5 show that our measure of the value-rent ratio for the housing stock varies responds much less to a change in z_t than the price-dividend ratio. The reason is that, while an increase in $\ln z_t$ increases discount rates, it also increases the expected growth rate of housing dividends, $\Delta \ln c_{t+1} - \Delta \ln z_{t+1}$. Since the consumption-housing ratio reverts to its mean, a high value today predicts an increase in housing expenditure in the future. This increases the current value of the housing stock, partly offsetting the increase in discount rates. In our model, houses are thus less risky than stocks and command a lower premium. The price dynamics also suggests the share of housing in total

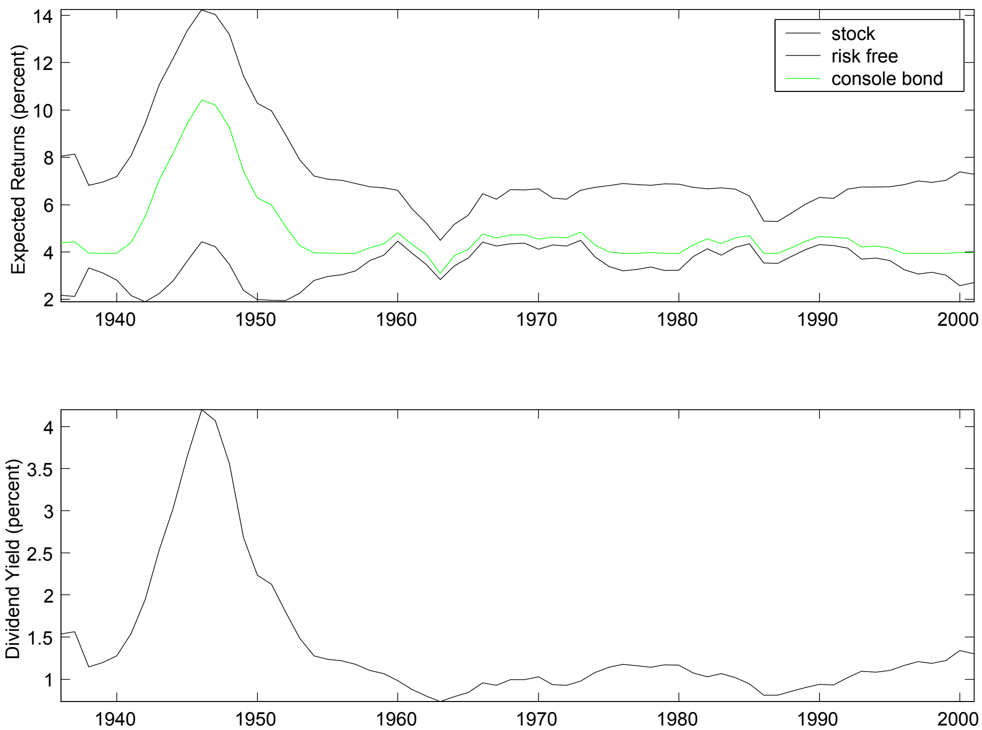

Figure 6: Expected returns and dividend-yield predicted by the model.

wealth, $c_t z_t \upsilon_t^s / (c_t \upsilon_t^s + c_t z_t \upsilon_t^s) = (1 - \alpha_t) \upsilon_t^h / (\upsilon_t^h + \upsilon_t^s)$ as another candidate variable for forecasting returns. We do not pursue this implication in the current paper, since it would require data on wealth.

Dependence on the Elasticities

Decreasing the elasticities of substitution ε and σ increases the impact of composition risk. Table 5 illustrates for the case of the intratemporal elasticity. As ε goes to one, the equity premium increases, the average riskless rate decreases and all asset prices become more volatile. A decrease in σ has a similar effect on premia and volatility, since it also increases $(\sigma - \varepsilon) / \sigma (\varepsilon - 1)$. However, it also has the standard effect of increasing the riskless rate since agents who like to substitute consumption push up the riskless rate in a growing economy. For this reason, we focus on the case of low risk aversion.

7 Time-series predictability

In this section, we explore predictability of returns in more detail. The results from the model of the previous section suggest that, in addition to the dividend yield, the non-housing share α_t should be a good forecasting variable. A higher expenditure share predicts high excess returns, similar to the dividend-yield. Panel A in Table 6 reports of the results of predictability regressions using simulated data from the model.

We can see that both the log dividend-yield and log alpha predict log excess returns with a positive sign. The R^2 using $\ln \alpha_t$ is also comparable to the R^2 from the log dividend-yield and rise with horizon. Panel B in Table 6 reports corresponding regression results with historical data over the whole sample, and over the postwar period. We can see that both the slope estimates and the R^2 are comparable to those in the model. In a horse race, the log expenditure share clearly outperforms the log dividend yield in the data, especially over longer forecasting horizons. In the model, the two variables are highly correlated, so we do not perform such a horse race.

Table 6: Time-series predictability of excess stock returns

Panel A: Model

Horizon	$\ln d_t/p_t^s$		$\ln \alpha_t$	
(years)	slope	R^2	slope	R^2
1	0.15	0.10	3.26	0.10
2	0.28	0.18	7.47	0.18
3	0.40	0.25	10.63	0.25
4	0.50	0.32	13.48	0.32
5	0.60	0.37	16.02	0.37

Panel B: Data

Horizon	$\ln d_t/p_t^s$			$\ln \alpha_t$			$\ln d_t/p_t^s$		$\ln \alpha_t$		
(years)	slope	t-stat	R^2	slope	t-stat	R^2	slope	t-stat	slope	t-stat	R^2
1	0.11	2.51	0.07	1.36	1.47	0.02	0.10	2.04	0.43	0.44	0.07
2	0.21	2.15	0.12	3.30	2.03	0.07	0.17	1.60	1.75	1.11	0.14
3	0.24	1.70	0.11	5.01	2.40	0.14	0.15	1.01	3.65	1.77	0.18
4	0.32	1.63	0.13	6.58	2.84	0.18	0.16	0.86	5.08	2.29	0.20
5	0.49	2.49	0.19	8.44	3.65	0.22	0.28	1.49	5.87	2.64	0.26
Postwar Sample											
1	0.11	2.25	0.08	1.42	1.68	0.03	0.10	1.84	0.50	0.13	0.08
2	0.21	1.95	0.12	3.68	2.24	0.08	0.16	1.39	2.14	0.75	0.14
3	0.24	1.53	0.11	6.25	3.21	0.20	0.10	0.66	5.30	2.11	0.21
4	0.32	1.41	0.11	8.63	3.95	0.28	0.06	0.30	8.09	3.04	0.28
5	0.52	2.26	0.17	10.73	4.92	0.30	0.15	0.81	9.24	3.43	0.31

NOTE: Panel A contains average slope and R^2 over 5000 simulated samples of length equal to 1936-2001. Panel B reports regressions using annual data during 1936-2001 and the 1947-2001 postwar period. T-statistics are based on Newey-West standard errors to correct for overlapping observations. Returns, dividend yields and alpha are all in logs.

8 Cross-section predictability

We now turn to the cross-sectional pricing implications of a model with housing. To investigate these implications, we use annual data on the 25 Fama-French portfolios. The 25 portfolios are the intersections of 5 portfolios formed on firm size and 5 portfolios formed on the ratio of book equity to market equity. For any candidate collection z of factors explaining the cross-section of excess returns, we can write $E(R^{ei}) = \beta^i \lambda$, where the betas β^i measure the covariance of excess returns R^{ei} with z. To estimate this specification, we use Fama-MacBeth regressions. The first step of this procedure regresses each excess return on z to obtain the betas. The second step estimates λ by regressing the cross section of average excess returns on β^i-estimates from the first step and a constant. Table 7 reports the results from these regressions.

Table 7: Cross-sectional evidence

	const.	R^m	SMB	HML			R^2
CAPM	5.71	5.12					0.08
	(1.30)	(1.03)					[0.04]
FF3	11.18	−2.75	2.24	6.58			0.86
	(2.48)	(−0.55)	(1.34)	(3.71)			[0.84]
		$\Delta \ln c_{t+1}$	$\Delta \ln \alpha_{t+1}$	$\Delta \ln c_{t+1} \cdot \ln \alpha_t$	$\Delta \ln \alpha_{t+1} \cdot \ln \alpha_t$	$\ln \alpha_t$	
CCAPM	11.09	0.02					0.58
	(3.95)	(3.79)					[0.56]
w/o housing	10.32	0.03					0.63
	(3.76)	(3.84)					[0.61]
CHCAPM	10.36	0.02	0.005				0.71
	(3.75)	(2.02)	(3.32)				[0.69]
scaled	8.08	0.02		0.00008		0.01	0.67
CCAPM	(3.85)	(3.20)		(0.28)		(1.93)	[0.63]
scaled	11.50	0.002	0.004	−0.0004	−0.00004	0.0008	0.82
CHCAPM	(4.89)	(0.025)	(3.15)	(−1.49)	(−0.09)	(0.15)	[0.78]

NOTE: Returns are on the 25 Fama-French portfolios on Ken French's website. Fama-MacBeth t-statistics are in round brackets. R^2 are in the last column, squared brackets indicated adjusted R^2. The sample period is 1936-2001.

Theory implies that the intercept in this cross-sectional regression should be zero. For all specifications in Table 7 except for the CAPM, the t-statistics of the intercept, which are the round brackets in the second column, are above 2. Thus, all of these specifications are rejected in the cross-section. However, the R^2 and t-statistics from these regressions are still interesting, because they may point to interesting factors. In terms of explaining the cross-section of returns, the CAPM does poorly over the sample period, with an R^2 of 8%. Although rejected, the CCAPM explains much more than the CAPM in terms of R^2, an impressive 58%. Interestingly, these results reverse when we include the Great Depression. During the sample period 1930-2001, the CAPM has an R^2 of 37%, while the CCAPM-R^2 drops to 10%. Even more interestingly, a CCAPM based on our measure of numeraire consumption (CCAPM w/o housing) seems to do somewhat better than the conventional measure. This finding suggests that our consumption definition may be interesting also for other consumption-based asset pricing applications which do not necessarily take into account housing.

Our model motivates a two factor model with α as the scaling variable. The two factors are the consumption growth rate and the growth rate of α. In fact, we can already do better than the standard CCAPM by scaling consumption growth with α. However, the growth rate of α is a significant factor. This consumption-housing CAPM, which we call CHCAPM, explains 71% of the cross-sectional variation in returns. The value of the regression coefficient on the growth rate of α, however, is not easily reconcilable with the parameter values we use for our time series results. Our results are related to those in Cochrane (1996) who documents that real investment growth helps pricing the cross section of returns. Moreover, Cochrane shows that the important component in real estate investment is residential real estate, not commercial real estate investment (Table 9 on page 615). His finding is consistent with our model, where residential real estate matters to consumers. Since movements in α captures risk at low frequencies, our results are also related to Parker and Julliard (2003). Their results show that low-frequency measures of consumption perform better in the CCAPM.

We can further improve by scaling both factors, consumption growth and α growth, by the level of α. This scaled CHCAPM model compares well with the Fama-French 3 factor model in terms of R^2. In addition to the market return, French and French (1997) use the return on a portfolio strategy which holds small firms in CRSP and shorts big ones (small-minus-big, SMB). Moreover, they use the return on a portfolio which holds high book-to-market firms and shorts those with low ones (high-minus-low, HML). Taken together, these three factors explains 86% of the cross section variation in excess returns, while the scaled CHCAPM explains 82%. The results for the scaled CHCAPM are even stronger when we include the Great Depression.

Figure 7 plots the fitted values from the various models against the realized values. Ideally, all observations in this figure should be along the 45^o line. Each observation is marked by two numbers. The first number stands for size ranging from the smallest, 1, to the largest, 5. The second number stands for the book-to-market category. Here, 1 stands for highest book-to-market, while 5 stands for lowest. From the left upper graph in the figure, we can see that fitted returns values on small value stocks are about the same as those for large growth stocks according to the CAPM. The resulting R^2 is therefore low. The right upper graph shows the success of the 3-factor Fama-French model, which predicts higher returns for small value stocks (labelled 15). The lower left graph is the CCAPM. We can see that the covariance of returns on small value stocks with consumption growth explains more of their relative performance than their covariance with the market. The lower right graph is the scaled CHCAPM, which compares well with the Fama-French model.

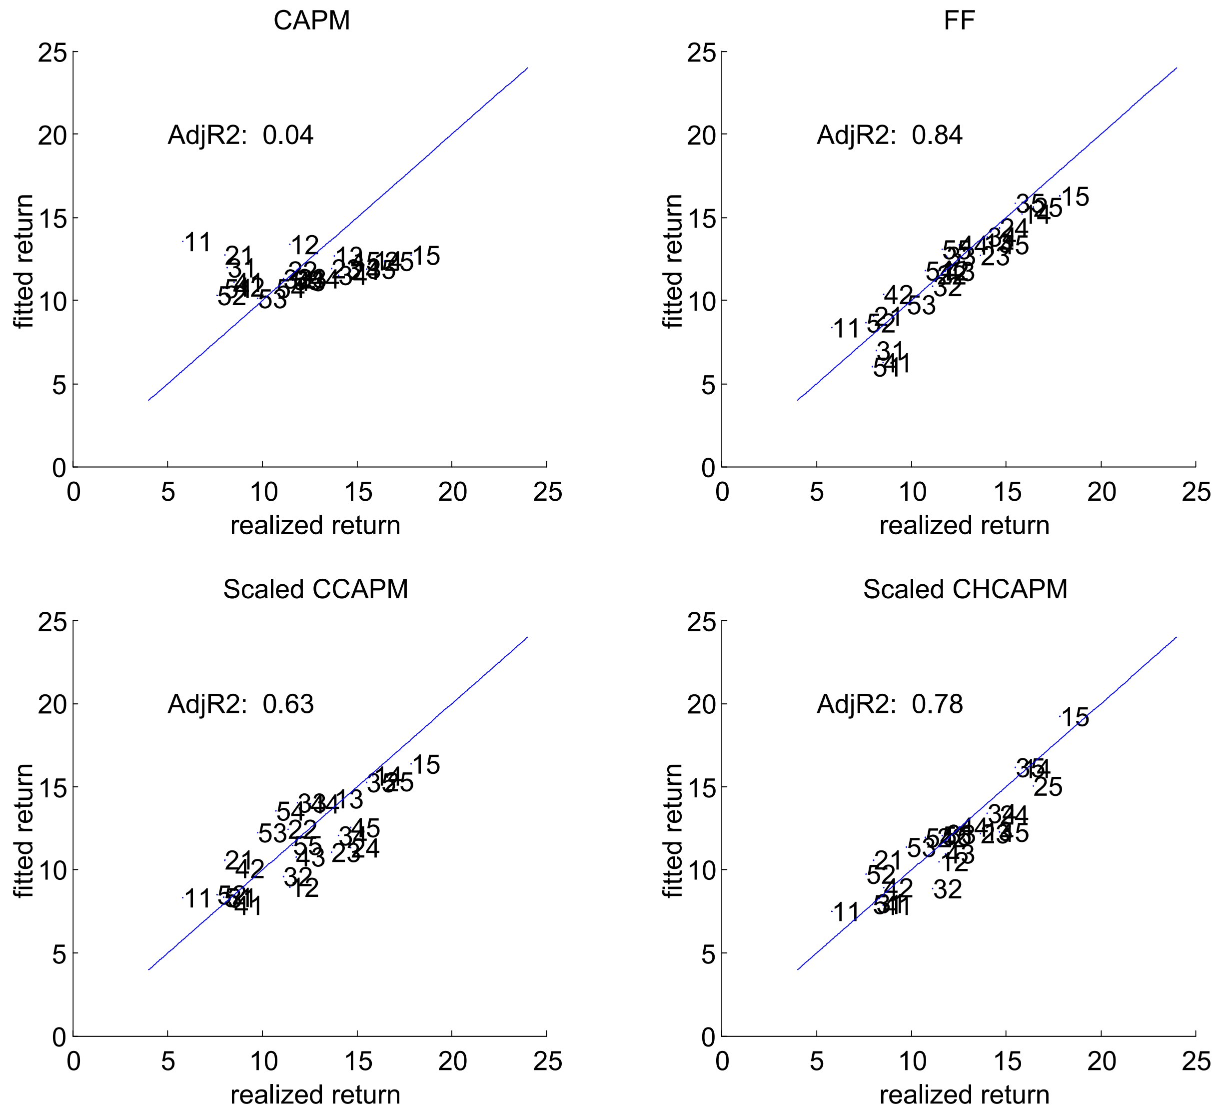

Figure 7: Fitted and realized returns on the 25 Fama-French portfolios based on various models.

9 Conclusion

We introduce an equilibrium model for asset pricing with housing. Agents in our model care about the composition of a consumption basket that contains shelter and other goods. Composition risk leads to higher equity and housing premia, higher stock return volatility, a lower riskfree rate which is not volatile, and lower bond premia. The model also implies that the dividend-yield and the non-housing expenditure share α forecast excess stock returns in the future. Moreover, the model motivates a two factor regression model with α as a scaling variable. The two factors are the consumption growth rate and the growth rate of α. Quantitatively, the model generates returns consistent with those in the data, when we calibrate the model to non-housing consumption data and housing expenditures. Remarkably, the expenditure share α predicts stock returns over time better than the dividend yield. Moreover, expenditure share data help explain the cross section of stock returns. This is interesting, because α is not based on asset market data. The calibration relies on a reasonable coefficient of relative risk aversion and values for the intratemporal elasticity of substitution that are in line with evidence in the literature. At these values, however, the model implies relative quantities and prices (rents) that are highly volatile. In particular, rent inflation is more volatile than CPI data on rents. More work with reliable relative quantity and price data is needed to accurately assess the role of composition risk.

References

[1] Abel, Andrew B. (1999). "Risk Premia and Term Premia in General Equilibrium." *Journal of Monetary Economics* 43, pp. 3-33.

[2] Attanasio, Orazio P. and Guglielmo Weber (1995). "Is consumption growth consistent with intertemporal optimization? Evidence from the Consumer Expenditure Survey." *Journal of Political Economy* 103, pp. 1121-57.

[3] Campbell, John Y. (1986). "Bond and Stock Returns in a Simple Exchange Model." *Quarterly Journal of Economics* 101, pp. 785-804.

[4] Campbell, John Y. (2001). "Consumption-based asset pricing." forthcoming in the *Handbook of the Economics of Finance*, George Constantinides, Milton Harris, and Rene Stulz eds., North-Holland, Amsterdam.

[5] Campbell, John Y. and John H. Cochrane (1999). "By force of habit: A consumption-based explanation of aggregate stock market behavior." *Journal of Political Economy*, 107, pp. 205-51.

[6] Caplin, Andrew, Sewin Chan, Charles Freeman, and Joseph Tracy (1997). *Housing Partnerships.* MIT Press, Cambridge, Massachusetts.

[7] Case, Karl E. and Robert J. Shiller (1989). "The efficiency of the market for single-family homes." *American Economic Review* 79, pp.125-37.

[8] Cocco, Joao (2000). "Hedging house price risk with incomplete markets." Working paper, London Business School.

[9] Cochrane, John H. (1991). "Production-based asset pricing and the link between stock returns and economic fluctuations." *Journal of Finance* 46, pp. 209-37.

[10] Cochrane, John H. (1996). "A cross-sectional test of an investment-based asset pricing model." *Journal of Political Economy* 104, pp. 572-621.

[11] Cochrane, John H. (1997). "Where is the market going? Uncertain facts and novel theories." *Economic Perspectives*, Federal Reserve of Chicago, 21, pp. 3-37.

[12] DeMarzo, Peter, Ron Kaniel and Ilan Kremer (2003). "Diversification as public good: Community effects in portfolio choice." Forthcoming *Journal of Finance.*

[13] Dunn, K. and Kenneth Singleton (1986). "Modeling the term structure of interest rates under non-separable utility and durability of goods." *Journal of Financial Economics* 17, pp. 27-55.

[14] Eichenbaum, Martin and Lars Peter Hansen (1990). "Estimating models with intertemporal substitution using aggregate time series data." *Journal of Business and Economic Statistics* 8, pp. 53-69.

[15] Eichenbaum, Martin, Lars Peter Hansen and Kenneth Singleton (1988). "A time series analysis of representative agent models of consumption and leisure choice under uncertainty." *Quarterly Journal of Economics* 103, 51-78.

[16] Fama, Eugene F. and Kenneth R. French (1997). "Size and book-to-market factors in earnings and returns." *Journal of Finance* 50, pp. 131-155.

[17] Flavin, Marjorie (2002). "Owner-occupied housing in the presence of adjustment costs: Implications for asset pricing and nondurable consumption." Working Paper, UC San Diego.

[18] Flavin, Marjorie and Takashi Yamashita (2002). "Owner-occupied housing and the composition of the household portfolio." *American Economic Review, March*, 345-62.

[19] Hanushek, Eric A. and John M. Quigley (1980). "What is the price elasticity of housing demand?" *Review of Economics and Statistics* 62(3), pp. 449-54.

[20] Heathcote, Jonathan and Morris Davis (2001). "Housing and the business cycle." Working paper, Georgetown University.

[21] Heaton, John (1993). "The interaction between time-nonseparable preferences and time aggregation." *Econometrica* 61, pp. 353-85.

[22] Heaton, John (1995). "An empirical investigation of asset pricing with temporally dependent preference specifications." *Econometrica* 63, pp. 681-718.

[23] Jagannathan, Ravi and Zhenyu Wang (1996). "The conditional CAPM and the cross-section of expected returns." *Journal of Finance* 51, pp. 3-53.

[24] Kullmann, Cornelia (2002). "Real estate and its role in asset pricing." Working paper, UBC.

[25] Lettau, Martin, and Sydney Ludvigson (2001a). "Consumption, aggregate wealth and expected stock returns," *Journal of Finance* 56, pp. 815-49.

[26] Lettau, Martin, and Sydney Ludvigson (2001b) "Resurrecting the (C)CAPM: A cross-sectional test when risk premia are time-varying." *Journal of Political Economy* 109, pp. 1238–1287.

[27] Lustig, Hanno (2001). "The Market Price of Aggregate Risk and the Wealth Distribution." Working paper, Chicago University.

[28] Lustig, Hanno and Stijn Van Nieuwerburgh (2002). "Housing Collateral, Consumption Insurance and Risk Premia." Working paper, Chicago University.

[29] Mayo, Stephen (1981). "Theory and estimation in the economics of housing demand." *Journal of Urban Economics* 10, pp. 95-116.

[30] Mehra, Rajnish and Edward Prescott (1985). "The equity premium puzzle." *Journal of Monetary Economics* 15, pp. 145-61.

[31] Ogaki, Masao and Carmen M. Reinhart (1998). "Measuring intertemporal substitution: The role of durable goods." *Journal of Political Economy* 106, pp. 1078-98.

[32] Ortalo-Magne, Francois and Sven Rady (1999). "Boom in, bust out: Young households and the house price cycle." *European Economic Review* 43, pp. 755-66.

[33] Parker, Jonathan and Christian Julliard (2003). "Consumption risk and cross-sectional returns." Working paper, Princeton.

[34] Prescott, Edward (1997). "On defining real consumption." *Review of Federal Reserve Bank of Minneapolis*, May/June, pp. 47-53.

[35] Runkle, David E. (1991). "Liquidity constraints and the permanent-income hypothesis: Evidence from Panel data." *Journal of Monetary Economics* 27, pp. 73-98.

[36] Santos, Tano and Pietro Veronesi (2001). "Labor Income and Asset Returns." Working paper, Chicago GSB.

[37] U.S. Department of Labor, "*Housing expenditures.*" Issues in Labor Statistics 02-02, 2002.

[38] Weil, Pierre (1989). "The equity premium puzzle and the riskfree rate puzzle." *Journal of Monetary Economics* 24, 401-21.

Appendix

This appendix compares our NIPA-based measures of house prices and returns with alternative measures. Davis and Heathcote (2001) use the price index for new residential investment from NIPA Table 7.6, line 38, as measure of house prices p_t^h. This series is a chain-type price index for investment in private residential structures starting in 1947, and it does not include the value of the land. This is the index that mimics our index the closest among all the indexes; the correlation of price changes between this index and our house price index is 0.80.

An alternative price index is provided by the Office of Federal Housing Enterprise Oversight (OFHEO). Starting in 1975, this index tracks the changes in the value of single family homes through repeat sales using the mortgage transaction data provided by Fannie Mae and Freddie Mac. The OFHEO index reflects the cost of structures and land, simultaneously controlling for the quality of the house. The series, however, does not go back very far. The correlation of price changes with our index is 0.71 over the 25 years where we have data on the OFHEO index.

The National Association of Realtors (NAR) publishes indexes that report median house prices starting early 1960s. The Bureau of the Census also reports median and average sale prices of houses sold in the United States since 1963. These indexes do not control for quality of the median house. The Census Bureau also publishes constant-quality price indexes that do not include the value of the land, but correct for the quality problem. These indexes are also available starting early 1960s.

Flavin and Yamashita (2001) use PSID data on house prices to estimate the housing returns over the 1968-1992 period. Unlike the other house price measures we just discussed, PSID house price data is at the homeowner level. Returns can therefore be computed for individual houses. There is, however, no rent data accompanying house prices in the PSID. Flavin and Yamashita therefore compute real housing returns as:

$$\frac{p_t^h + \overline{r_f}\ p_{t-1}^h + \tau \text{Propertytax}_t}{p_{t-1}^h} = \frac{p_t^h}{p_{t-1}^h} + \overline{r_f} + 0.33 \times 0.025 \tag{17}$$

where $\overline{r_f}$ denotes the average real short term interest rate, the personal tax rate τ is set to be 33% and property tax rate is set to 2.5%. Flavin and Yamashita set $\overline{r_f}$ to be 5% which seems too high in our sample. We compute $\overline{r_f}$ using our data.

Equation (14) explains how we compute the real return on housing. There are two main reasons that make our house price measure superior to other alternatives in our analysis. First, ours is the only house price measure that goes back until the 1930s. Second, we have rents (housing expenditures) that correspond to our house price series. Table 8 reports summary statistics on individual housing returns from Flavin and Yamashita (2001, Table 1A) and our aggregate housing returns series. We compute aggregate housing returns using Flavin and Yamashita's (FY) return definition (17), and using our definition based on rent data (14).

Table 8: Various measures of real returns on housing

	FY definition (Eq.17) 1968-1992		Our definition (Eq.14) 1968-1992
	PSID data	Our data	Our data
mean	6.59	1.80	1.97
std	14.24	2.74	2.81

NOTE: This table reports mean and standard deviation of real housing returns. The first column reports the findings in Flavin and Yamashita (2001, Table 1A), the second column is (17) evaluated with our house price index. The third column is (14) evaluated with our house price index. Returns are deflated using the price index for nondurables and services, excluding shoes & clothing, and housing services.

Table 8 shows that average returns on individual housing are more than 3 times as high as those on aggregate housing. The difference in standard deviations is even more striking. Returns on individual houses are more than 5 times as volatile as returns on the U.S. housing stock as a whole. The last columns in Table 8 shows that rent data only matters little for the volatility of aggregate returns.

Table 9 presents real returns on housing using the FY definition of returns with different house price indexes that are discussed before. The mean returns on housing are around 2-3% for all indexes and time periods, and the standard deviation of returns are in the 1.5-3% range except the last column. In the last column, housing return statistics are calculated for each state separately using OFHEO state level house price indexes and then averaged. Going from the aggregate to the state level, the volatility of housing returns almost doubles. Idiosyncratic housing returns are still more than 2 times as volatile as state-level housing returns.

Table 9: Real returns on housing using the FY definition

	Our data	DH data	OFHEO data	
			aggregate	state level
	1947-2000	1948-2000	1975-2000	1975-2000
mean	1.96	2.00	2.82	2.51
std	2.21	1.70	3.19	5.86

NOTE: This table reports mean and standard deviation of real housing returns using the FY definition (17). The first column is based on our house price index. The second column is based on the price index for new residential investment as used in Davis and Heathcote (2001). The third and fourth columns are based on the OFHEO price indexes at aggregate and state levels, respectively. Returns are deflated using the price index for nondurables and services, excluding shoes & clothing, and housing services.