

Expected Returns and the Business Cycle: Heterogeneous Goods and Time-Varying Risk Aversion

Lars A. Lochstoer, Columbia University and London Business School*

First draft: May 2004

November 3, 2008

Abstract

We propose a representative agent habit formation model where preferences are defined over both luxury goods and basic goods. The model matches the equity risk premium, risk free rate, and volatilities. From the intratemporal first order condition we can substitute out basic good consumption and the habit level, yielding a stochastic discount factor driven by two *observable* risk factors: luxury good consumption growth, and the relative price of the two goods. We estimate these processes and find them to be heteroskedastic, implying time-variation in the conditional volatility of our stochastic discount factor. These dynamics occur both at the business cycle frequency and at a lower, "generational frequency". Consistent with the model's predictions, we empirically document the existence of these two frequencies in the equity premium, where the lower frequency is quantitatively more important. We also investigate the model's implications for bond returns, and relate their predictability to the low frequency variation in the conditional volatility of the relative price growth.

*This paper is a substantially revised version of Chapter 1 of my dissertation at the Haas School of Business, University of California at Berkeley. I thank Pierre Collin-Dufresne, Roger Craine, Tom Davidoff, Francisco Gomes, Keener Hughen, Georg Kaltenbrunner, Shimon Kogan, Jacob Sagi, Richard Stanton, Motohiro Yogo and especially my adviser Greg Duffee and an anonymous referee for helpful comments. I thank seminar participants at Amsterdam University, Berkeley, Boston College, Carnegie-Mellon, CEPR Gerzensee 2005, Columbia, London Business School, Norwegian School of Economics, Oslo School of Management (BI), Stanford, University of British Columbia, University of Southern California, University of Washington at Seattle, and the WFA 2005. I thank the Dean Witter Foundation and Senter for Pengepolitisk og Finansiell Forskning for financial support. All errors are mine alone. Contact info: Lars Lochstoer, Columbia University, Uris Hall 405B, 3022 Broadway, New York, NY 10010. E-mail: LL2609@columbia.edu

1 Introduction

Consumption-based asset-pricing models tie the dynamic behavior of the stock market risk premium to the real economy. In order to fit the unconditional moments of aggregate consumption growth and asset returns, these models generally predict a counter-cyclical risk premium (e.g., Chan and Kogan, 2002; Campbell and Cochrane, 1999; Melino and Yang, 2003). On an intuitive level, the cycle that one refers to is the business cycle. However, the equity premium generated by these models is typically extremely slow-moving; it essentially follows a "generational" cycle too long to reflect a pure (NBER) business cycle phenomenon.¹ Empirically, there is little evidence so far of a tight link between the state of the economy (the business cycle) and measures of the equity risk premium, although these are usually correlated (e.g., Fama, 1991; Fama and French, 1989; Lettau and Ludvigson, 2001). In particular, the predictability of excess equity market returns appears to be strong at forecasting horizons much longer than the length of a typical business cycle, and some of the most successful forecasting variables, like the price-dividend ratio or the short-term interest rate, are much more persistent than typical business cycle variables.² If the risk factor that generates this predictability is not a pure business cycle phenomenon, what is its relation to macro economic conditions? Furthermore, is there a relation between the equity risk premium and the business cycle, independent of this slower moving factor?

In this paper, we provide answers to these questions from the perspective of a representative agent external habit formation model (e.g., Abel, 1990; Campbell and Cochrane, 1999; Wachter, 2006). Different from what is usual in habit formation models, the habit is in our model effectively observable, which enables the habit parameters to be *estimated* using available macro data, instead of relying on a calibration to asset pricing moments. Following Ait-Sahalia, Parker, and Yogo (2004) we assume that a representative agent has non-homothetic preferences over basic goods and luxury goods, where basic good consumption is subject to an external minimum consumption level, while luxury good consumption is not. This utility specification captures the notion that luxury good consumption growth

¹I use "generational" as a term for cycles that are substantially longer than standard business cycles, e.g. 15-20 year cycles, or with a persistence corresponding to the post-war persistence of the aggregate price/dividend ratio.

²See Boudoukh, Michaely, Richardson, and Roberts (2007) for recent evidence on the predictive power of the aggregate dividend yield.

is a better proxy for agents' marginal utility than total consumption growth because it is a better measure of consumption investors can adjust at the margin.³ Importantly, using the intratemporal first order condition of the agent, we can substitute out the habit using the relative price between luxury good and basic good consumption, and so the marginal utility of the agent, and thus the pricing kernel, is in effect observable. In other words, luxury good consumption and the relative price contain information that is sufficient to describe the dynamics of the habit.

We construct an aggregate measure of luxury good consumption from data in the detailed consumption tables published by the Bureau of Economic Analysis, available from 1959Q1 to 2006Q4. The remaining goods in the aggregate consumption basket are deemed as basic, and we use the implicit price deflators of each consumption good to construct the aggregate relative price. We show that this measure of luxury good consumption growth is more volatile than aggregate consumption growth and more highly correlated with stock returns, consistent with the findings in Ait-Sahalia, Parker, and Yogo (2004). These are important features of the data in terms of the model's ability to match the equity premium puzzle (Mehra and Prescott, 1985).

To evaluate the model's asset pricing implications, we estimate a joint Constant Correlation EGARCH-in-mean (Nelson, 1991; Bollerslev, 1990) specification for luxury good consumption and relative price growth. This specification is relatively parsimonious, while allowing for time-variation in both the first and second moment of each series. We find large and significant time-variation in the conditional volatility of the two series. In particular, the conditional volatility of luxury good consumption growth operates at a business cycle frequency and is high recessions and low in expansions. The conditional volatility of the relative price growth, however, operates at a lower, generational frequency and is positively related to the aggregate equity market dividend yield and a measure of the aggregate consumption-wealth ratio (the *CAY*-variable of Lettau and Ludvigson, 2001). Thus, the model predicts time-variation in the conditional volatility of the stochastic discount factor at *both* the business cycle and a generational frequency. This time-variation is *not* uncovered by relying on asset price dynamics, but estimated from real macro data. The calibrated model predicts that most of the fluctuations in equity and bond risk premiums are related to the lower frequency movements in macro economic risk, as captured by the conditional volatility of

³A related empirical fact is that the consumption of the rich, who are more likely to consume luxury goods, is more volatile than and more highly correlated with stock returns than that of the poor (see, e.g., Mankiw and Zeldes, 1991; Vissing-Jorgensen, 2001).

the relative price growth. In other words, the slow-moving time-variation in the conditional volatility of surplus consumption, which is the fundamental driver of time-varying volatility of the pricing kernel in habit formation models (see Campbell and Cochrane, 1999), show up in the macro data as time-varying conditional volatility in the relative price between consumption baskets that are more (basic goods) or less (luxury goods) exposed to the habit. When basic good consumption is low (high) relative to the habit level, the relative price is more (less) sensitive to shocks to the ratio of luxury good to basic good consumption. We show that the estimated conditional volatility of the relative price growth predicts excess stock and bond returns.

The model further predicts that the business cycle variation in the amount of risk mainly shows up in Sharpe ratios. This is consistent with the findings in Lettau and Ludvigson (2007), who show that the Sharpe ratio is a business cycle variable that is high in recession and low in expansions. The equity risk premium is also higher in recessions than expansions in the model, but these fluctuations are relatively hard to pick up in forecasting regressions with the sample sizes we have available. Nevertheless, we provide empirical evidence of a business cycle component in the equity risk premium that is separate from the already identified lower frequency component. Finally, we show that excess returns to nominal risk-free bonds are also time-varying. In particular, the Cochrane and Piazzesi (2006) factor predicts bond returns at all maturities in the model, consistent with the data. Thus, the expectations hypothesis does not hold in the model (Fama and Bliss, 1986; Campbell and Shiller, 1991; Cochrane and Piazzesi, 2006).

In terms of the unconditional moments, the model is able to account for the equity premium puzzle, the risk-free rate puzzle, and the excess volatility puzzle with a relative risk aversion over luxury good gambles of 10. This is consistent with the findings in Ait-Sahalia, Parker, and Yogo (2004), who use a different data set. Further, the model matches the unconditional risk premiums on nominal risk-free bonds, but predicts too high volatility of medium term bond returns.

The model in this paper is complementary to Bansal and Yaron (2004) and Lettau and Wachter (2007) who argue that slow-moving volatility in *aggregate* consumption growth causes the ultra-low frequency movements in the risk premium. Lettau and Ludvigson (2007), however, argue that estimated consumption growth volatility cannot explain conditional asset pricing puzzles. In this paper, we do not take a stance on the process for aggregate consumption growth, but instead focus on the implications of a two-good model. The model in this paper is thus related to the asset pricing literature on heterogenous goods models

(see, e.g., Lustig and Van Nieuwerburgh, 2005; Piazzesi, Schneider and Tuzel, 2006; Yogo, 2006). Further, different persistence of fundamental shocks is also studied in Calvet and Fisher (2005) who investigate the impact of shocks with different persistence levels on asset prices.

The paper proceeds as follows. Section 2 presents the model. Section 3 gives the data and details of its construction. In section 4, we estimate and calibrate the model parameters. Section 5 gives the implications of the model for relevant conditional and unconditional asset pricing moments, while section 6 concludes.

2 The Model

The model in this paper is a consumption-based representative agent, heterogeneous goods model in the spirit of Ait-Sahalia, Parker, and Yogo (2004). In particular, the representative household is assumed to have Cobb-Douglas preferences over basic and luxury good consumption,

$$U(L_t, B_t; X_t) = E_t \sum_{j=1}^{\infty} \beta^j \frac{[L_{t+j}^\alpha (B_{t+j} - X_{t+j})^{1-\alpha}]^{1-\gamma}}{1-\gamma} \quad (1)$$

where $\gamma > 0$ is the curvature coefficient, β is the time discounting parameter, X_t is an external difference habit and α is the share of total surplus consumption ($C_t - X_t$) that goes to luxury good consumption. Note that if $\alpha = 0$ the model collapses to the one-good, external habit formation model of Campbell and Cochrane (1999). If $\alpha = 1$, the model is similar to the model in Ait-Sahalia, Parker and Yogo (2004), who assume that the utility function is separable in the two goods. The habit level in the model can be thought of as a minimum basic good consumption level related to the standard of living. For instance, U.S. household real, per capita expenditures on gas and electricity has more than doubled from 1959 to 2006. However, few would today consider themselves well-off for affording electricity to operate the current standard of household appliances, but instead rather poor if one could not afford it. Thus, such an increase in real consumption is unlikely to be associated with a substantial utility gain, as a habit over the standard of living has developed.⁴

⁴The existence of a minimal level of basic goods consumption can also be motivated by both a subsistence level and the presence of *commitment* goods. Chetty and Szeidl (2003) argue that agents pre-commit to consumption levels of certain goods, e.g. flow of housing consumption, cellular phone service, gym-membership, etc. Such commitments mean shocks to wealth will primarily be reflected in non-commitment goods that are easily adjusted at the margin. We generally think of commitment levels as endogenous, while X is assumed to be external, but Chetty and Szeidl provide a theoretical justification for this modeling

The utility function is non-homothetic since the agent must consume the externally given level of basic goods $X > 0$. Thus, the wealthier an investor is relative to the habit level, the larger the ratio of her luxury good consumption to basic good consumption.⁵ The distinction between luxury goods and basic goods is motivated by Ait-Sahalia, Parker, and Yogo (2004) who show empirically that luxury good consumption (goods typically consumed only by the rich) do better at explaining asset prices than aggregate consumption. The above utility function achieves this in a general equilibrium setting. In this model, aggregate luxury good consumption growth is a better proxy for agents' marginal utility than total consumption growth because it is a better measure of consumption investors can adjust at the margin. Both aggregate basic good and total consumption growth are contaminated by the unobservable consumption floor, X .

Let the basic good be the numeraire good and P_t denote the relative price of luxury goods. The intratemporal first order condition of the household is then

$$\frac{P_t L_t}{B_t - X_t} = \frac{\alpha}{1 - \alpha}, \quad (2)$$

and the total expenditure is

$$C_t = P_t L_t + B_t. \quad (3)$$

Combining these two equations, and ignoring proportional constant terms involving α , we can express the agent's per period utility function in terms of total consumption and the relative price:⁶

$$u(C_t; P_t, X_t) = \frac{[(C_t - X_t) P_t^{-\alpha}]^{1-\gamma}}{1 - \gamma}. \quad (4)$$

choice as a reduced form representation. Another interpretation is that luxury goods are a euphemism for frictionless goods consumed mainly by the wealthy (stockholders).

⁵In standard Cobb-Douglas preferences, the proportion of income spent on each good is equal to its share in the utility function; α and $(1 - \alpha)$. This is not the case here due to the presence of the habit level. In fact, from the intratemporal first order conditions, the equilibrium luxury good expenditure share is given by

$$\frac{P_t L_t^*}{C_t} = \frac{\alpha(C_t - X_t)}{C_t},$$

which will be counter-cyclical given the slow-moving nature of a habit level.

⁶ $P_t L_t = \frac{\alpha}{1-\alpha} (B_t - X_t) \iff P_t L_t = \frac{\alpha}{1-\alpha} (C_t - X_t - P_t L_t) \iff P_t L_t = \alpha (C_t - X_t)$. We are ignoring constant terms involving α in equations (4), (5), and (6) for clearer exposition.

The marginal utility of total consumption is then

$$MU_t = (C_t - X_t)^{-\gamma} P_t^{\alpha(\gamma-1)} \quad (5)$$

$$= L_t^{-\gamma} P_t^{-\gamma+\alpha(\gamma-1)}. \quad (6)$$

The habit level is unobserved, but in the two-good setting of this paper the marginal utility can be expressed as a function of only observable variables. In particular, as the second line in the above expression for marginal utility shows, we can substitute out the habit level as a function of the relative price and luxury good consumption using equations (2) and (3). In other words, the relative price contains information that in conjunction with luxury good consumption is sufficient to make the habit in effect observable. This is very useful as the calibration and estimation of the model now can be done relying on observable variables.

The stochastic discount factor is then

$$M_t = \beta \exp(-\gamma \Delta l_t - \delta \Delta p_t), \quad (7)$$

where $\delta \equiv \gamma - \alpha(\gamma - 1)$.

2.1 State Processes

There are two state processes which dynamic behavior need to be determined in this model; luxury good consumption growth (Δl_t) and the price of luxury good consumption relative to basic good consumption (Δp_t). We specify these processes with the following in mind. First, to allow for time-variation in the amount of risk we must allow for heteroskedasticity in the joint dynamics of Δl_t and Δp_t . Second, due to the limited amount of data available, we need the specification to be as parsimonious as possible. Given these consideration, we let these variables follow a constant correlation EGARCH-in-mean process (Bollerslev, 1990; Nelson, 1991). In particular,

$$\Delta l_t = a_l + b_l \sigma_{l,t}^2 + \sigma_{l,t} \varepsilon_{l,t} \quad (8)$$

$$\log \sigma_{l,t}^2 = \omega_l + \beta_{1,l} |\varepsilon_{l,t-1}| + \beta_{2,l} \varepsilon_{l,t-1} + \beta_{3,l} \log \sigma_{l,t-1}^2 \quad (9)$$

and

$$\Delta p_t = a_p + b_p \sigma_{p,t}^2 + \sigma_{p,t} \varepsilon_{p,t} \quad (10)$$

$$\log \sigma_{p,t}^2 = \omega_p + \beta_{1,p} |\varepsilon_{p,t-1}| + \beta_{2,p} \varepsilon_{p,t-1} + \beta_{3,p} \log \sigma_{p,t-1}^2 \quad (11)$$

where $\text{corr}(\varepsilon_{l,t}, \varepsilon_{p,t}) = \rho_{lp}$ and $\varepsilon_{l,t}, \varepsilon_{p,t} \sim N(0, 1)$. From an economic standpoint it is important to include the volatility in the mean specification for the variables. Equilibrium considerations often lead to a trade-off between the volatility and the mean in the optimal consumption of a good.⁷ In our case, the relative price is directly related to the unobservable habit level through surplus consumption, $C_t - X_t$. In habit models, such as Campbell and Cochrane (1999), the expected conditional growth rate of surplus consumption is high exactly when its conditional volatility is also high. The EGARCH process further accommodates asymmetric response in conditional volatility to positive or negative shocks. This allows us to interpret the volatility as counter- or pro-cyclical and is again consistent with the habit literature where a negative shock to surplus consumption increases its conditional volatility. Finally, the EGARCH specification ensures that volatility can never become negative while at the same time allowing for persistence and cyclical in volatility.

Under these assumptions there are two state-variables in the real economy, $\sigma_{l,t}$ and $\sigma_{p,t}$. Since the shocks are assumed to be normally distributed, the one-period log risk-free rate is

$$r_{f,t} = -\ln E_t[M_{t+1}], \quad (12)$$

$$= (-\ln \beta + \gamma a_l + \delta a_p) + \gamma (b_l - \gamma/2) \sigma_{l,t}^2 + \delta (b_p - \delta/2) \sigma_{p,t}^2 - \gamma \delta \rho \sigma_{l,t} \sigma_{p,t}, \quad (13)$$

and the maximal conditional Sharpe ratio in the economy approximately satisfies⁸

$$SR_{\max} \equiv \max_i \frac{E_t(r_{i,t+1}) - r_{f,t+1} + \frac{1}{2} \sigma_{i,t}^2}{\sigma_{i,t}} \lesssim (\gamma^2 \sigma_{l,t}^2 + \delta^2 \sigma_{p,t}^2 + 2\gamma \delta \rho \sigma_{l,t} \sigma_{p,t})^{1/2}. \quad (14)$$

Thus, time-variation in both the risk-free rate and the maximum Sharpe ratio are driven by time-varying volatility in innovations to the relative price and luxury good consumption. Since these variables are observable, we can estimate their mean and volatility dynamics.

⁷Said differently: the optimal amount of intertemporal substitution undertaken in general depends on the volatility of the consumption process.

⁸This approximation comes from assuming that returns, luxury good consumption growth, and the relative price growth are jointly, conditionally log-normally distributed. See Campbell (1999).

Cash Flows.

We will be concerned with the value of nominal bonds and the claim to aggregate dividends (the market portfolio). To value these we must specify their real cash flows. For nominal bonds, this amounts to specifying the process for inflation. Note that the nominal log risk-free rate is given by

$$r_{f,t}^{nom} = -\ln E_t \left[M_{t+1} \frac{\Pi_t}{\Pi_{t+1}} \right], \quad (15)$$

where Π_t is the price level at time t . If inflation is high, a nominal one-period default-free bond pays off less than if inflation is low. The nominal stochastic discount factor is thus

$$M_{t+1}^{nom} = M_{t+1} \frac{\Pi_t}{\Pi_{t+1}}. \quad (16)$$

Inflation can be written

$$\Delta\pi_{t+1} = E_t [\Delta\pi_{t+1}] + \sigma_\pi \varepsilon_{\pi,t+1}, \quad (17)$$

where $\Delta\pi_{t+1}$ is the log inflation rate, $\varepsilon_{\pi,t} \sim N(0, 1)$, $E[\varepsilon_{l,t}\varepsilon_{\pi,t}] = \rho_{l\pi}$ and $E[\varepsilon_{p,t}\varepsilon_{\pi,t}] = \rho_{p\pi}$; that is, we allow inflation to be correlated with the real risk factors in the economy. Thus, nominal bonds will in general not have the same risk premium dynamics as otherwise similar real bonds due to an indirect inflation risk premium. We write "indirect" as inflation itself is not a source of risk in the economy. Nominal default-free bonds simply have risky payoffs from the perspective of the representative agent, who cares about real consumption. Inflation is highly persistent, and following Wachter (2006) we let expected inflation follow an AR(1) process:

$$E_t [\Delta\pi_{t+1}] = c_\pi + \phi E_{t-1} [\Delta\pi_t] + \sigma_{E\pi} \sigma_\pi \varepsilon_{\pi,t}. \quad (18)$$

Note that innovations to expected inflation are assumed to be perfectly correlated with innovations to realized inflation. This allows us to write realized inflation as:

$$\Delta\pi_{t+1} = c_\pi + \phi E_{t-1} [\Delta\pi_t] + \sigma_{E\pi} \varepsilon_{\pi,t} + \sigma_\pi \varepsilon_{\pi,t+1}, \quad (19)$$

$$= c_\pi + \phi \Delta\pi_t + \theta \sigma_\pi \varepsilon_{\pi,t} + \sigma_\pi \varepsilon_{\pi,t+1}, \quad (20)$$

where $\theta = \sigma_{E\pi} - \phi$. Thus, realized inflation follows an ARMA(1,1) process, which can be easily estimated from the data via maximum likelihood. This specification is motivated partly empirically, as previous research has documented an important MA(1) component in realized inflation (e.g., Wachter, 2006), and partly by parsimony as inflation as specified here

only introduces one additional state variable, $E_t[\Delta\pi_{t+1}]$, which evolution parameters can be easily estimated from the data.

Finally, we assume that aggregate log real dividend growth is given by:

$$\Delta d_{t+1} = c_d + (\delta_1 + \delta_2 \sigma_{p,t}) \varepsilon_{d,t+1}, \quad (21)$$

where $\varepsilon_{d,t+1} \sim N(0, 1)$, $E[\varepsilon_{l,t} \varepsilon_{d,t}] = \rho_{ld}$ and $E[\varepsilon_{p,t} \varepsilon_{d,t}] = \rho_{pd}$. Thus, aggregate dividend growth is unpredictable, as in, e.g., Campbell and Cochrane (1999) and Wachter (2006), but with time-varying conditional volatility, as in, e.g., Bansal and Yaron (2004). The latter feature is empirically motivated, as we find that the conditional volatility of the relative price growth predicts realized dividend growth volatility in the data.

3 Data

The financial data on aggregate market returns, Treasury bills, and Treasury bonds used in this study are downloaded from the CRSP database. The sample period is 1959Q2 - 2006Q4, of reasons explained below.

Consumption Data. The model calls for nondurable luxury good consumption data. However, there is no agreed upon source of aggregate luxury good consumption data. What is more, a relatively long data series is necessary for the purpose of estimating the above Constant Correlation EGARCH-in-mean model with its 13 parameters. Ait-Sahalia, Parker, and Yogo (2004), which is the benchmark study for the use of luxury good consumption for asset pricing purposes, use as their main series sales data from some luxury retail companies (e.g., Tiffany's) to construct an aggregate data series. They rely on this series as opposed to consumption data from the Bureau of Economic Analysis to avoid significant components of basic goods in the luxury good consumption measure. Unfortunately, this data cannot be used directly for the model in this paper. First, this data is annual from 1960 - 2001. While the data clearly constitutes luxury good and not basic good consumption, 41 observations is insufficient to reliably estimate the above EGARCH system. In addition, the series ought to be of a quarterly frequency or higher to capture volatility dynamics well. These authors also provide year-on-year quarterly data from 1986 - 2001, which is a slightly longer, although still quite short, sample in terms of number of observations (60 vs. 41). However, the data should also span a relatively long time period in order to capture any low frequency fluctuations in

the amount of risk. Finally, the mentioned data series consist mainly of durable goods as they cover jewelry expenditures. In particular, the relative price series they employ is the relative price of retail jewelry stores. The model presented in this paper, however, specifies nondurable luxury good consumption.

Given these data issues, we therefore construct an alternative measure of nondurable luxury good consumption and relative prices based on the detailed NIPA tables from the Bureau of Economic Analysis. This data is quarterly and goes back to 1959Q1. In particular, I classify sub-categories of aggregate nondurable and services expenditures as luxury good consumption. I apply the following three criteria when picking sub-categories.

1. The good is not predominantly a necessary good.

This rules out food and clothing, for instance. While high-end restaurant visits and fashion certainly should be classified as luxury good consumption, the NIPA data unfortunately does not allow us to separate such consumption from the more "basic" visits to, e.g., fast-food restaurants and generic brand clothes stores and most of the food and clothing consumption is of the latter flavor.

2. The good is not predominantly a basic good; it is likely to be consumed relatively more by the wealthy

This rules out many forms of discretionary spending on recreational activities, such as movie theater tickets. While going to the movies is clearly not a necessary good, it should be considered a basic good. It is quite inexpensive, and it does not separate the richer from the poorer as movie theater visits are not scalable in quality to any significant extent. For instance, as you get more wealthy, you cannot improve your movie experience to a significant extent by paying more. The ability to increase the quality of the good consumed by spending more money is an important feature of luxury goods and discussed next.

3. Non-satiability.

Luxury goods are scalable in quality. For instance, hotel visits and travel are clearly discretionary. While many of us may be relatively quickly satiated in terms of the *amount of days* we travel per year, there are few limits (satiation) on the *amount you spend* per hotel visit and the flight. By choosing increasingly luxurious travel arrangements (business/first class airfare and luxury hotels) instead of low budget travel arrangements, the luxury good aspect of the good can be increased.

Table 1 - Summary Statistics

Table 1: This table reports summary statistics for log luxury good consumption growth, Δl , and log relative price growth, Δp , series. The table includes relevant moments for aggregate nondurable, per capita log consumption growth, as well as log excess equity market returns (CRSP). Correlation both with returns reported as contemporaneous, r_1^e , and returns reported a quarter earlier, r_2^e , are given due to the effects of time-aggregation in the consumption series (see, e.g., Campbell, 1999). All moments are based on annual growth rates and returns. The sample is 1960 - 2006.

	<i>Mean</i>	<i>St.Dev</i>	<i>Correlations</i>			
			Δl	Δp	Δc	$\Delta l + \Delta p$
Δl	2.37%	4.89%	1			
Δp	0.27%	2.32%	-0.02	1		
Δc	2.25%	1.39%	0.62	0.17	1	
r_1^e	5.88%	15.77%	0.15	0.30	0.25	0.27
r_2^e	5.86%	18.73%	0.36	0.38	0.33	0.49

Applying these criteria leaves us with 4 sub-categories of US domestic consumption of nondurables and services: Hotel stays, Air travel, Beauty Parlors and Health Clubs, and Private Flying. While these categories do have basic good components (perhaps with the exception of Private Flying), they are likely to be consumed more as one becomes more wealthy. We use this basket of goods as a proxy for aggregate nondurable and services luxury good consumption and construct the luxury good consumption series as a value-weighted index of the underlying consumption series. I construct the implicit price deflator for both luxury goods and basic goods from the underlying implicit price deflator series provided in the NIPA tables. The Appendix gives details of the data construction. This basket of goods is not as cleanly identified as luxury good consumption as the data in Ait-Sahalia, Parker, and Yogo (2004), but the availability of data and the consistency in terms of the basket of goods versus the construction of the relative price it allows are very important for the estimation of the model. Note that the measure of luxury good consumption used in this paper has similar cyclical properties as the measure constructed by Ait-Sahalia, Parker, and Yogo (2004) and the correlation between the two measures is 0.5. Finally, contamination from basic goods in the luxury good basket will work *against* the model and thus the classification is conservative.

Table 1 gives relevant summary statistics for the luxury good consumption and relative price data used in this study. Real luxury good consumption growth is 3.5 times as volatile as real aggregate consumption growth. The relative price growth is also more volatile than ag-

gregate consumption growth, and all three series have about the same correlation with excess equity market returns. The latter is given both as the contemporaneous annual correlation (r_1^e) and as annual consumption and price growth leading annual returns by one quarter (r_2^e). The leading correlations are given due to the time aggregation that is implicit in measured consumption and prices as given by NIPA (see, e.g., Campbell, 1999). The final column gives the correlation of real luxury good expenditure growth with excess equity returns (i.e., $\Delta l_t + \Delta p_t$). With a maximum correlation with equity returns of 0.49, this measure is more highly correlated with excess equity market returns than aggregate consumption is. Figure 1 shows the sample of quarterly real, per capita luxury good consumption growth and price growth, versus aggregate consumption growth (dashed line) and NBER recession indicators (bars). The figure confirms visually the higher volatility of the two series compared with aggregate consumption growth. Further, luxury good consumption is clearly pro-cyclical, with low growth in recessions. The relative price growth also tends to be lower in recessions, but the relationship with business cycles is not nearly as strong. For instance, in the recession in 1980, the relative price in fact increased, while both aggregate consumption and luxury good consumption growth decreased.

So, how does this series compare to the luxury good series given in Ait-Sahalia, Parker, and Yogo (2004)? Figure 2 shows that the two series move together, as one would expect if they to some extent capture luxury good spending. The correlation between the annual growth rates of these two series in the available overlapping sample (1961-2001) is 0.49.⁹ The correlation between the Ait-Sahalia, Parker, and Yogo (2004) luxury good consumption series and aggregate returns is 0.26 (contemporaneous) and 0.42 (consumption leading returns by 1 quarter), versus 0.27 and 0.49 for the measure used in this paper. It is clear from the figure that the measure used in Ait-Sahalia, Parker, and Yogo (2004) is more volatile than the measure used in this paper. This can be partly attributed to the fact that the basket of goods used to construct luxury good consumption in this paper is aggregate and span different classes of consumption goods which are not perfectly correlated, while the core data series in Ait-Sahalia et al. is from a subset of luxury good retail stores that mainly specialize in jewelry sales.

Inflation Data. The inflation data used in this study is the implicit price deflator of basic goods. This price deflator is constructed using the quarterly price series of aggregate

⁹The correlation between aggregate consumption growth and the APY (2004) luxury growth series is lower at 0.37.

Figure 1

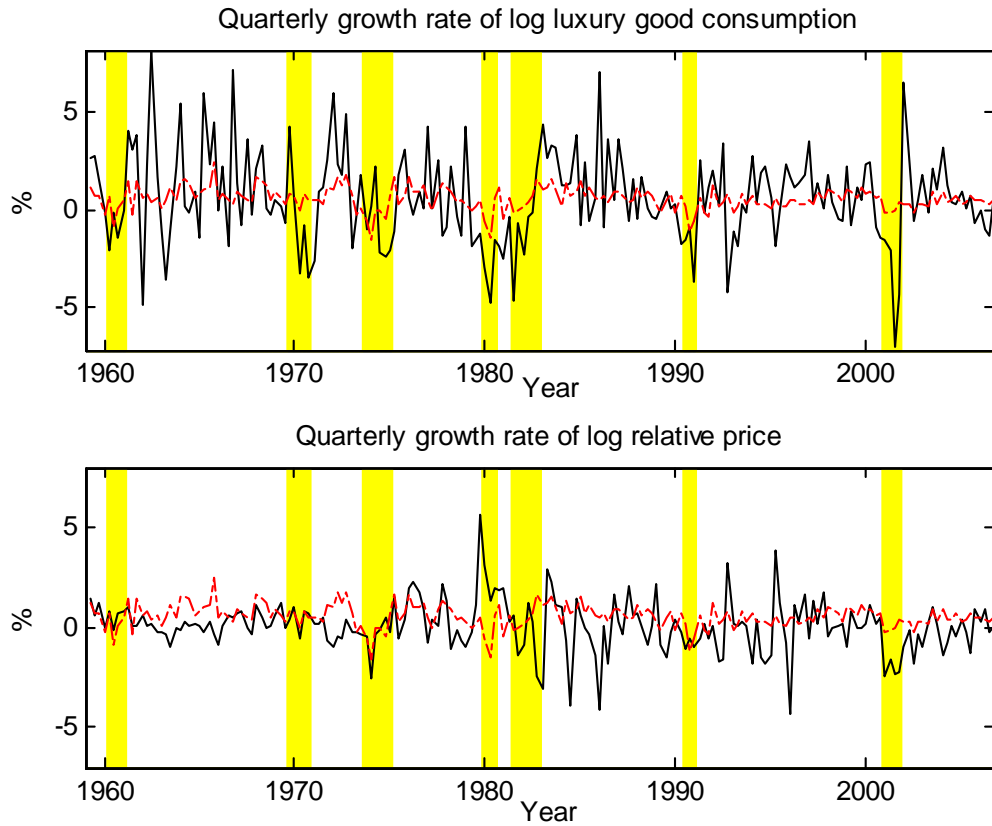


Figure 1: The upper figure shows quarterly real log luxury good consumption growth versus aggregate consumption growth (dashed line) and NBER recessions (bars). The lower figure shows the relative price growth versus aggregate consumption growth (dashed line) and NBER recessions (bars). The sample is 1959Q2 - 2006Q4.

Figure 2

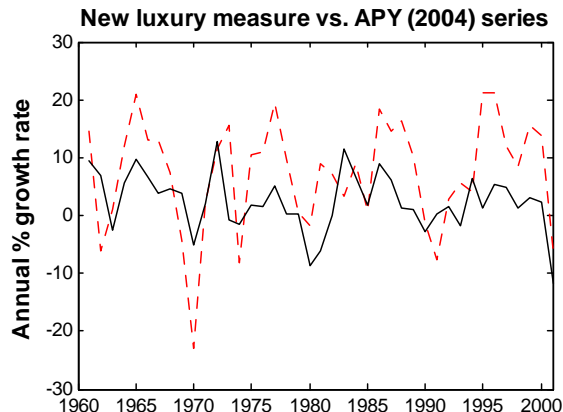


Figure 2: The figure shows the annual log luxury good consumption growth series given in Ait-Sahalia, Parker, and Yogo (2004), dashed line, versus the annual log luxury good consumption growth of the series constructed in this paper, solid line. The correlation between the two series is 0.49. The sample is 1961 - 2001.

nondurable and services consumption and the constructed price deflator for luxury good consumption. The construction methodology is explained in detail in the Appendix. Since luxury good consumption is a relatively small fraction of aggregate consumption, the resulting inflation series is empirically very close to the inflation series obtained using only the price deflators of aggregate nondurable and services consumption. The annual inflation rate over the sample is 3.61% with an annualized standard deviation of 1.31%.

Dividend Data. Aggregate stock market dividends are obtained from Boudoukh, Michaely, Richardson, and Roberts (2007), who construct an aggregate price-dividend ratio that takes into account share repurchases in addition to cash dividends using quarterly data from CRSP and Compustat. This data is annual and we use the sample 1959 - 2006 to construct sample moments and for estimation of the parameters of the dividend growth specification in equation (21). Log annual real dividend growth in the sample has a mean of 2.58% with a standard deviation of 12.07%. The average level of the price dividend ratio is 27.28, the volatility of the log price-dividend ratio is 0.274, and the annual autocorrelation of the log price-dividend ratio is 0.868.

4 Estimation and Calibration

4.1 Estimation

The joint EGARCH-in-mean specification given in equations (8) to (11) above is estimated by maximum likelihood. Details about the estimation procedure are given in the Appendix.

Before presenting the results, there are two issues concerning time-aggregation of the macro data that needs to be addressed. The model is solved and estimated at the quarterly frequency, in line with the availability of the data. While the model assumes that consumption takes place at the end of each quarter, this is not the case in the data. The consumption and price data are averages over the quarter. This induces spurious autocorrelation and smooths the series. Working (1960) shows that even though a variable (say, consumption growth) is actually i.i.d., time-aggregation induces a first-order autocorrelation of 0.25 and reduces the variance by a factor of 3/2. First, any dynamic behavior in the mean should not be due solely to time-aggregation. We therefore allow for an AR(1) term in the mean in the empirical specification to control for the induced temporal dependency. Following Wachter (2006), the AR(1) term will *not* feature in the specification used for the calibrated model. Second, when calibrating the model, the volatility of the series is important for matching relevant moments and thus for the chosen preference parameters. We discuss how we deal with the effect of time-aggregation on volatility in the section below on model calibration.

Panel A of Table 2 shows the parameter estimates from the bivariate, constant correlation EGARCH-in-mean model. Focusing first on the results for luxury good consumption growth, the volatility parameters are all significant at the 10% level or more. The ARCH term, $\beta_{1,l}$, equals 0.41 and the GARCH term, $\beta_{3,l}$, equals 0.77, which implies positive persistence in volatility. Further, the asymmetric response term, $\beta_{2,l}$, is -0.17 , which means that a negative shock to luxury good consumption growth increases its volatility. The upper half of Figure 3 shows the estimated volatility (dashed line) over the sample. The volatility operates at a business cycle frequency and tends to be high in recessions and low in expansions. For the relative price growth, the ARCH term, $\beta_{1,p}$, equals 0.29 and the GARCH term, $\beta_{3,p}$, equals 0.94, which implies positive and high persistence in volatility. Again, the asymmetric response term, $\beta_{2,p}$, is -0.12 is negative, which means that a negative shock to the relative price increases its volatility. The lower half of Figure 3 shows the estimated volatility (dashed line) over the sample. The volatility is much more persistent than that

Table 2 - Estimated Parameters

Table 2: Panel A reports maximum likelihood estimates of the parameters of the bivariate constant correlation EGARCH-in-mean process for luxury good consumption growth and relative price growth. The sample period is 1959Q3 - 2006Q4; 190 quarterly observations. Panel B reports a restricted version of the EGARCH model. Panel C reports parameter estimates of the processes assumed for inflation and the aggregate stock market dividends. Standard errors are reported in paranthesis.

Panel A: EGARCH estimation					
	Δl_t -series	(s.e.)		Δp_t -series	(s.e.)
a_l	0.0096**	(0.0041)	a_p	0.0001	(0.0013)
b_l	-11.5888	(11.3236)	b_p	4.7261	(13.7792)
$ar(1)$	0.1012	(0.0922)	$ar(1)$	0.3033**	(0.0953)
ω_l	-2.1351	(1.4717)	ω_p	-0.7399*	(0.4097)
$\beta_{1,l}$	0.4135**	(0.1637)	$\beta_{1,p}$	0.2855**	(0.0843)
$\beta_{2,l}$	-0.1743*	(0.0941)	$\beta_{2,p}$	-0.1186	(0.0750)
$\beta_{3,l}$	0.7727**	(0.1788)	$\beta_{3,p}$	0.9432**	(0.0404)
<i>Likelihood value</i>		1116.55	ρ	-0.2189**	(0.0778)
Panel B: EGARCH estimation - Restricted Model					
	Δl_t -series	(s.e.)		Δp_t -series	(s.e.)
a_l	0.0015	(n/a)	a_p	-0.0001	(0.0007)
b_l	5.6048	(n/a)	b_p	8.9974	(n/a)
$ar(1)$	0.2222**	(0.0643)	$ar(1)$	0.3183**	(0.0875)
ω_l	-1.8465	(1.2558)	ω_p	-0.8307**	(0.3761)
$\beta_{1,l}$	0.3228**	(0.1273)	$\beta_{1,p}$	0.2886**	(0.0861)
$\beta_{2,l}$	-0.1460*	(0.0799)	$\beta_{2,p}$	-0.1219*	(0.0696)
$\beta_{3,l}$	0.7980**	(0.1514)	$\beta_{3,p}$	0.9340**	(0.0368)
<i>Likelihood value</i>		1114.53	ρ	-0.2188**	(0.0765)
<i>LR-test</i> = -2(1114.53-1116.55) = 4.04 $\sim \chi^2(3)$. <i>p-value</i> = 0.26					
Panel C: Inflation and Dividends					
	Inflation, π_t	(s.e.)	Moments of annual dividends, Δd_A		
c_π	0.0096**	(0.0023)	c_d	0.0258	(0.0176)
ϕ	0.9282**	(0.0317)	$\tilde{\delta}_1$	0.0439	(0.0279)
θ	-0.3647**	(0.0809)	$\tilde{\delta}_2$	4.7807**	(2.4193)
σ_π	0.0036**	(0.0007)	$\tilde{\rho}_{ld}$	0.44	
$\rho_{l\pi}$	0.0553	(0.0728)	$\tilde{\rho}_{pd}$	0.51	
$\rho_{p\pi}$	-0.1392*	(0.0722)			
R^2	0.70				

for luxury good consumption growth; it operates at a lower, "generational" frequency rather than at a business cycle frequency.

For both of the variables, the coefficient on the variance in the mean specification is insignificant. Luxury good consumption growth tends in the sample to be low if the conditional volatility is high, while the opposite is true for the relative price growth. The relevant parameters b_l and b_p are very important for the dynamic behavior of the real risk-free rate (see equation (13)), and therefore for bond risk premiums. In the main part of the paper, we will consider a version of the EGARCH process that restricts b_l and b_p to values that maximizes the model's fit to the real risk-free rate. Panel B of Table 2 shows parameter estimates from this restricted version of the joint EGARCH process. The values chosen for b_l and b_p will be discussed in more detail in the Calibration section, so for now simply note that the restricted model is not rejected in favor of the unrestricted model in a Likelihood Ratio test (see bottom of Table 2). Further, the restrictions have only very small impact on the parameters governing the volatility processes. This can be seen visually in Figure 3. The dashed lines show the original volatility estimates, while the solid lines show the volatility estimates from the restricted model. For the luxury good consumption growth, the correlation between the two is 0.964, while for the relative price growth the correlation is 1.000. Finally, both the AR(1) terms are relatively small, and we cannot reject a first-order autocorrelation of 0.25, which is a benchmark number for autocorrelation induced by time-aggregation (Working, 1960).¹⁰

Inflation. Inflation in the model is the dollar price change of the basic good.¹¹ Panel C of Table 2 shows the parameters that obtain from estimating the ARMA(1,1) process for inflation given in equation (20) on the available data. As expected, inflation is very persistent, with a quarterly autocorrelation parameter of 0.93. The MA(1) term is negative and significant, and we can from its value back out the relative magnitude of the shock to expected inflation: $\sigma_{E\pi} = \theta - \phi = -0.3647 + 0.9282 = 0.5635$. In words, shocks to expected inflation are perfectly positively correlated with, but smaller than, shocks to realized inflation in the estimated specification. Finally, innovations in inflation have low correlation with innovations to the real risk factors. Thus, inflation risk premia are economically small in

¹⁰The standard errors of the coefficient on the variance in the mean specification increase slightly if the AR(1) term is left out of the estimation, but the point estimates do not change significantly.

¹¹This notion of inflation is different from the implicit price deflator of nondurable and services consumption used in standard one-good models. Empirically, however, the two measures are very close with a time-series correlation of 0.99.

Figure 3

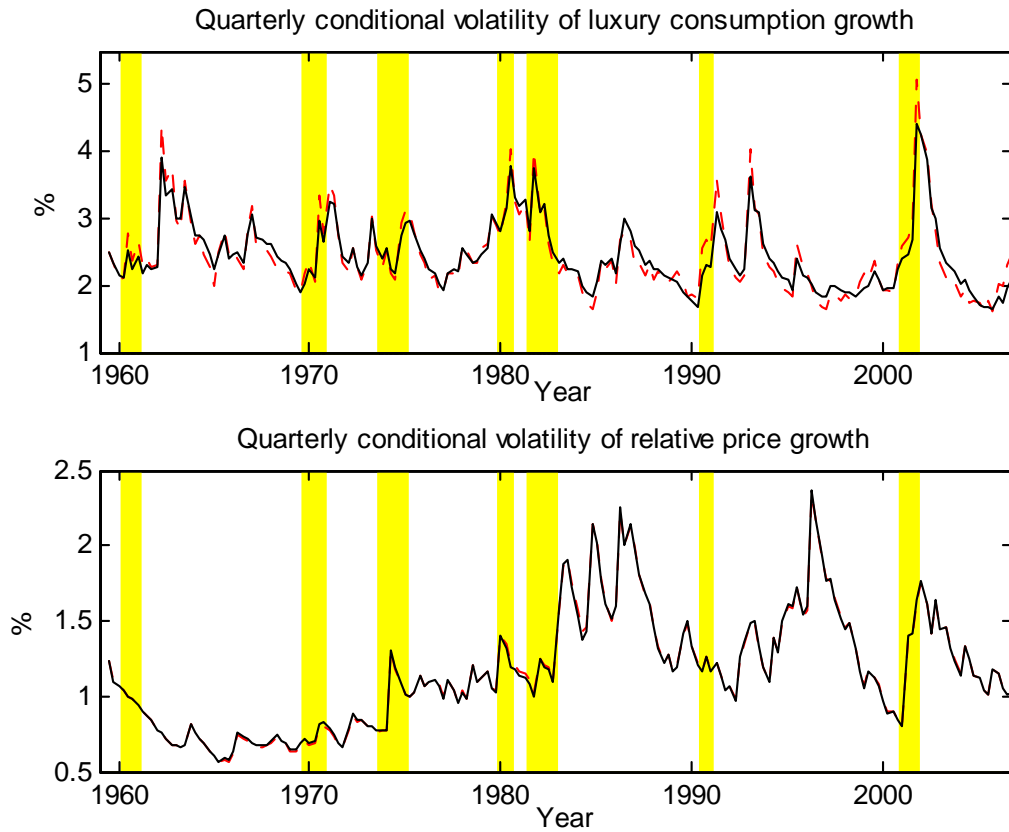


Figure 3: The upper plot shows the estimated conditional volatility of quarterly real log luxury good consumption growth. The dashed line shows the estimates from the unrestricted model, while the solid line shows the estimates from the restricted model. The lower plot shows the estimated conditional volatility of quarterly the relative price growth. The dashed line (mainly hidden behind the solid line) shows the estimates from the unrestricted model, while the solid line shows the estimates from the restricted model. The bars are NBER recession indicators. The sample is 1959Q3 - 2006Q4.

this model and so the dynamic behavior of nominal bond risk premia are predominantly inherited from the dynamic behavior of real bond risk premia.

Dividends. Panel C of Table 2 shows the result of estimating the dividend process in equation (21) using the annual dividend growth data. The parameters that determine the conditional volatility, δ_1 and δ_2 , are estimated by first running the regression

$$|\Delta d_t - \overline{\Delta d}| = \tilde{\delta}_1 + \tilde{\delta}_2 \hat{\sigma}_{p,t-4}^2 + \eta_t \quad (22)$$

on the annual dividend growth data, where η_t is the usual error term. Note that the conditional volatility of the relative price growth is lagged a full year so it is a valid predictive variable. The model calls for a quarterly regression, so the above regression is misspecified. We deal with this explicitly in the calibration section. For now, note that $\tilde{\delta}_2$ is positive and statistically significant. The conditional volatility of luxury good consumption growth is not included in this specification as it empirically does not show up as significant (not reported).

The contemporaneous correlation between annual aggregate dividend growth and annual luxury good consumption and relative price growth are only 0.08 and 0.17, respectively. These numbers are low, especially considering that real annual aggregate earnings growth (obtained from Robert Shiller's web site) has a correlation of 0.49 with luxury good consumption growth and 0.25 with relative price growth. This disparity in the correlations can, however, be reconciled by considering lead and lag correlations. For instance, one year lagged luxury good consumption growth have a correlation of 0.27 with the subsequent year's dividend growth, while annual real aggregate earnings growth has a correlation of 0.41 with the subsequent year's dividend growth. Such predictability is perhaps not surprising given the discretionary aspect of stock dividends. Time-averaging and reporting delays in consumption and price data means dividends may also lead these variables. To resolve these data issues, while keeping a parsimonious specification for dividends (which keeps the number of state variables low), I use a measure of correlation that takes lead and lag effects into account.

In particular, define $\tilde{\rho}_{ld} = \frac{\sum_{j=-1}^1 \text{cov}(\Delta l_{t-j}, \Delta d_t)}{\sigma(\Delta l_t)\sigma(\Delta d_t)} = 0.44$ and $\tilde{\rho}_{pd} = \frac{\sum_{j=-1}^1 \text{cov}(\Delta p_{t-j}, \Delta d_t)}{\sigma(\Delta p_t)\sigma(\Delta d_t)} = 0.51$.¹² These correlations are more in line with those we find with respect to earnings growth and, reassuringly, they make the model implied correlation of excess stock returns with luxury

¹²Using the same method for the calculation of the correlation between annual real aggregate consumption growth and dividend growth yields gives a correlation coefficient of 0.39; less than that for any of the two "new" factors.

good consumption and relative price growth close to those in the data (see Section 5). Panel C of Table 2 summarizes the annual moments that are used as a basis for the calibration of the dividend process in the model.

4.2 Calibration

The model is calibrated at a quarterly frequency, in line with the available data.

The Preference Parameters. There are three preference parameters that must be calibrated - the time-discounting parameter β , the luxury good consumption share α , and the curvature parameter γ .

First, consider the luxury share α . In equilibrium, we have that

$$\alpha = \frac{P_t L_t}{C_t} \frac{C_t}{C_t - X_t}. \quad (23)$$

That is, α is the luxury good expenditure share times the inverse of the expenditure share of total consumption in excess of the minimum basic good consumption level (X_t). The latter quantity is the inverse of an aggregate "surplus consumption ratio" ($\frac{C_t - X_t}{C_t}$; see, e.g., Campbell and Cochrane, 1999). Since the habit is unobservable, this variable is also in principle unobservable. However, it is possible to indirectly put some bounds on its average level by looking to both the habit formation literature and the literature on consumption commitments. From the habit formation literature, Campbell and Cochrane (1999) calibrates the average surplus consumption ratio using aggregate asset pricing moments to be 5.7%. Given that the sample average of $P_t L_t / C_t$ is 2%, the implied luxury share is $\alpha = 0.02 / 0.057 = 0.35$. Chetty and Szeidl (2007), on the other hand, study consumption commitments and find using U.S. consumption data that the average level of committed consumption expenditures is about 50% – 65% of total consumption.¹³ This implies a surplus consumption ratio of 35% – 50%, and we get $\alpha = 0.04 - 0.06 \approx 0.05$. Since a low α increases the importance of the relative price (δ increases) and since asset pricing dynamics at two frequencies is at the core of the estimated model, I choose $\alpha = 0.05$ in the benchmark calibration. Admittedly, this way of determining α is not very accurate; the studies cited are one-good models and the parameter's bounds are not very tight. On the other hand, it does not rely directly on the asset pricing moments the calibrated model will be compared to.

¹³Chetty and Szeidl (2005) show how the presence of consumption commitments provide a micro foundation for an aggregate habit level.

Next, consider the curvature parameter, γ . This parameter is not the same as the relative risk aversion of the agent ($RRA = \frac{-WV_{WW}}{V_W}$). The presence of a habit level confounds the usual interpretation of γ , as discussed in, e.g., Campbell and Cochrane (1999). In the model in this paper, however, per period utility can be written as $u_t \propto L_t^{1-\gamma} P_t^{1-\gamma+\alpha(\gamma-1)}$ by using equations (2), (3) and (4). In this case, $\frac{-L_t u_{LL}}{u_L} = \gamma$, and so we can view γ as the relative risk aversion coefficient over atemporal luxury good consumption gambles. Since the rationale for considering luxury good consumption in part is to offer an explanation of the equity premium puzzle, we set $\gamma = 10$, which is the maximum value of risk aversion suggested by Mehra and Prescott (1985). This is also the target range for the risk aversion parameter in the empirical tests in Ait-Sahalia, Parker, and Yogo (2004). We choose the maximum value of this parameter to maximize the possible Sharpe ratio in the economy. Given these values for α and γ , we have that $\delta = \gamma - \alpha(\gamma - 1) = 9.55$.

Finally, we restrict the time-discounting parameter, β , to be 0.9999. This value is less than 1, which is an economic upper bound for this variable, while its high value helps the model match the average level of the real risk-free rate. A similar value for this parameter has been used by, among others, Jermann (1998) and Boldrin, Christiano and Fisher (2000). A high value for β decreases the unconditional level of risk-free rate, as can be seen in equation (13).

The State Processes Revisited. Before we solve for asset prices in the model, it is useful to revisit the expression for the real risk-free rate

$$r_{f,t} = -\ln \beta + \gamma a_l + \delta a_p + \gamma (b_l - \gamma/2) \sigma_{l,t}^2 + \delta (b_p - \delta/2) \sigma_{p,t}^2 - \gamma \delta \rho \sigma_{l,t} \sigma_{p,t}. \quad (24)$$

With the parameter estimates of the joint EGARCH-in-mean process in hand, we are in a position to test the implied relation between the real risk-free rate and the conditional volatilities of luxury good consumption growth and the relative price growth empirically. This is a useful exercise at this point both to ensure that the estimated conditional volatilities in fact are related to the real risk-free rate and to help further with the calibration of the model. In particular, the model has rich implication for conditional asset pricing moments, and it is therefore important for the interpretation of the model to verify that its first-order implications for the dynamic behavior of the risk-free rate are in line with the data. Consider the regression:

$$r_{f,t} = \beta_0 + \underset{(-)}{\beta_1} \hat{\sigma}_{l,t}^2 + \underset{(-/+)}{\beta_2} \hat{\sigma}_{p,t}^2 + \underset{(-)}{\beta_3} \hat{\rho} \hat{\sigma}_{l,t} \hat{\sigma}_{p,t} + \varepsilon_t, \quad (25)$$

where the model predicted signs of the regression coefficients are given in parenthesis below the coefficients. The negative sign predicted for $\beta_1 = \gamma (b_l - \gamma/2)$ comes about due to the fact that $\hat{b}_l < 0$ and $\gamma > 0$. The sign on $\beta_2 = \delta (b_p - \delta/2)$ restricts the value of δ relative to \hat{b}_p , while the sign of $\beta_3 = -\delta\gamma$ must be less than zero as $\gamma, \delta > 0$. Running the regression in (25) using the quarterly sample (1959Q3-2006Q4) gives:

$$r_{f,t} = \beta_0 + \underset{(3.5209)}{6.0478\hat{\sigma}_{l,t}^2} + \underset{(15.1902)}{40.1584\hat{\sigma}_{p,t}^2} - \underset{(70.63)}{115.5323\hat{\rho}\hat{\sigma}_{l,t}\hat{\sigma}_{p,t}} + \varepsilon_t, \quad (26)$$

where White (heteroskedasticity corrected) standard errors are given in parenthesis. The first and the last coefficients are significant at the 10% level, while the second coefficient is significant at the 5% level. The fact that $\hat{\beta}_1$ is greater than 0 is, given that the point estimate of \hat{b}_l is -11.59 , a rejection of the model. However, the \hat{b}_l parameter is estimated with considerable uncertainty; its 95% confidence bounds are $(-34.23, 11.05)$. The sensitivity of the risk-free rate in the model to the estimated \hat{b}_l and \hat{b}_p is problematic given the large standard errors on these parameter estimates. For instance, a different sign of \hat{b}_l and/or \hat{b}_p leads to opposite risk-premium dynamics as bonds become positively or negatively correlated with the state-variables.¹⁴

Given these concerns, we let b_l and b_p be calibrated to $\hat{\beta}_1$ and $\hat{\beta}_2$ from the above regression. Using the calibrated values if γ and α , this gives $\hat{b}_l = 5.6048$ and $\hat{b}_p = 8.9974$. Further, the high value of γ and the high average sample growth rate of luxury good consumption (2.37% p.a.), causes a risk-free rate puzzle. With only 47 years of data, however, this average growth rate is also estimated with a large standard error. We therefore, restrict the mean for luxury good consumption growth to equal 1.4% p.a. which allows the model to match the average level of the risk-free rate. This implies a restriction on a_l . We then re-estimate the joint EGARCH process for Δl_t and Δp_t imposing the calibrated values of a_l , b_l and b_p and verify through a likelihood ratio test that the restricted model cannot be rejected in favor of the unrestricted model (see Panel B of Table 2).¹⁵ The p -value of the test of the restricted versus the unrestricted model is 0.26. Thus, the assumed parameters of the state processes

¹⁴A more technical point is simply the fact that bond prices do not converge if b_l is not sufficiently greater than zero for reasonable preference parameters. This happens because bonds in this case are assets that pay off increasing amounts in (not as fast) increasingly bad states.

¹⁵Strictly speaking, the calibration of b_l and b_p should be implemented as an iterative process as the estimates of the conditional volatilities depend on the values for b_l and b_p . Empirically, however, the conditional volatility series from the restricted model are very close to those of the unrestricted model, so iterating in this fashion is not quantitatively important. For simplicity and clarity of exposition, I opt for a simpler, more transparent calibration strategy.

Table 3 - Benchmark Calibration

Table 3: This table collects the parameter values used for the benchmark calibration of the model.

Δl_t -parameters		Δp_t -parameters		Preference-parameters		π_t -parameters		Δd_t -parameters	
a_l	0.0004	a_p	-0.0006	β	0.9999	c_π	0.0006	c_d	0.0065
b_l	5.6048	b_p	8.9974	γ	10	ϕ	0.9282	δ_1	0.0155
ω_l	-1.7756	ω_p	-0.8307	α	0.05	θ	-0.3647	δ_2	3.95
$\beta_{1,l}$	0.3228	$\beta_{1,p}$	0.2886			σ_π	0.0036	ρ_{ld}	0.44
$\beta_{2,l}$	-0.1460	$\beta_{2,p}$	-0.1219			$\rho_{l\pi}$	0.0553	ρ_{pd}	0.51
$\beta_{3,l}$	0.7980	$\beta_{3,p}$	0.9340			$\rho_{p\pi}$	-0.1392		

used when analyzing the model's implications are consistent with the available data. The strategy of matching the EGARCH model's parameters to the time-series behavior of the risk-free rate is similar in spirit to term structure models using observable factors (e.g., Ang and Piazzesi, 2003, Duffee, 2006). These models also match the dynamic behavior of the risk-free rate as well as possible using observable factors. Note that we have not used any conditional information to match the dynamic behavior of any asset's risk *premium* in the calibration. Thus, the model implied risk premium dynamics are "out-of-sample" from the perspective of the calibration strategy.

Finally, due to issues related to time aggregation as discussed in an earlier section, the $AR(1)$ terms in the mean specification of Δl_t and Δp_t are ignored in the model calibration, and ω_l and ω_p are set so as to match the annual, *unconditional* sample moments of $\sigma(\Delta l_t)$ and $\sigma(\Delta p_t)$, respectively, while a_p is set to match $E[\Delta p]$.¹⁶ Table 3 summarizes the parameter values used in the benchmark calibration, and Panel A of Table 4, show relevant unconditional moments of luxury good consumption and relative price growth of the calibrated model.

Inflation. The inflation process is calibrated using the estimated parameters of the ARMA(1,1). Since the data is quarterly, the correlations of the shock to inflation and the

¹⁶The annual variances are less smoothed by the time-aggregation as the annual growth numbers are based on year on year Q4 growth rates. In fact, in both cases the variance of annual growth rates are close to $4 * \frac{3}{2}$ of the variance of quarterly growth rates, consistent with the results in Working (1960). A similar strategy is adopted in Ait-Sahalia, Parker, and Yogo (2004) who scale the covariance between the consumption variables and returns by a factor of 2 to account for time-aggregation in their empirical tests.

shocks to luxury good consumption growth and relative price growth are easy to compute and we use the sample values of these.

Dividends. The dividend growth in the model is of a quarterly frequency, while the estimation in the previous section was based on annual dividend growth. To parameterize the quarterly dividend growth specification needed for the model (see equation (21)), we run a simple simulated method of moments exercise. In particular, we simulate data from the model, calibrated at a quarterly frequency where dividends are determined as in equation (21). Given an initial guess of δ_1 and δ_2 , we then simulate 1,000 samples of the same length as in the data and run the *annual* regression as in equation (22). We take the average regression coefficient from the 1,000 samples and compare these to $\tilde{\delta}_1$ and $\tilde{\delta}_2$ as estimated in the data. This guides a new guess of the true δ_1 and δ_2 , and we iterate on this procedure until the model generated $\tilde{\delta}_1$ and $\tilde{\delta}_2$ matches the $\tilde{\delta}_1$ and $\tilde{\delta}_2$ we find in the data. Table 3 shows the resulting parameter values. In the calibrated model, the conditional volatility of quarterly dividend growth lies between 2.7% (5th-percentile) and 7.0% (95th-percentile). These fluctuations in cash flow risk contributes to the time-variation in the equity risk premium at the same, generational frequency as $\sigma_{p,t}^2$. Finally, we also in the same procedure find the quarterly conditional correlations between dividend growth and the risk factors that correspond to the annual estimates given in the previous section. Since the conditional correlations between quarterly dividend growth and the two risk factors - luxury good consumption growth and relative price growth - are positive, shocks to realized dividends carry a positive risk premium.

5 Model Results

5.1 Model Solution

There are three state-variables in the model: the conditional volatility of luxury good consumption growth, $\sigma_{l,t}^2$, the conditional volatility of the relative price growth, $\sigma_{p,t}^2$, and the expected inflation rate, $E_t[\Delta\pi_{t+1}]$. All asset prices can be found as functions of these variables by numerically integrating the price functions. Below are definitions of key variables in the model.

The price-dividend ratio of the aggregate equity claim is given by

$$\frac{P_t^E}{D_t} = E_t \left[\sum_{j=1}^{\infty} M_{t,t+j} \frac{D_{t+j}}{D_t} \right], \quad (27)$$

where D_{t+j} is the real dividend payment at time $t + j$. Note that the price-dividend ratio can be expressed as a function only of the real state variables, $\sigma_{l,t}^2$ and $\sigma_{p,t}^2$, since dividends are given in real terms as in equation (21).¹⁷ The nominal return to equity is

$$R_{t+1}^E = \frac{P_{t+1}^E + D_{t+1}}{P_{t+1}^E} \frac{\Pi_{t+1}}{\Pi_t}. \quad (28)$$

The price of a nominal zero-coupon bond that pays \$1 in N periods is

$$P_{N,t}^{\$} = E_t \left[M_{t,t+N} \frac{\Pi_t}{\Pi_{t+N}} \right]. \quad (29)$$

The return on the bond is then

$$R_{N,t+1}^{\$} = \frac{P_{N-1,t+1}^{\$}}{P_{N,t}^{\$}}. \quad (30)$$

The continuously compounded yield to maturity on the bond is given by

$$y_{N,t}^{\$} = -\frac{1}{N} \ln P_{N,t}^{\$}. \quad (31)$$

The nominal risk-free rate is just the one-period nominal zero-coupon bond:

$$r_{f,t}^{\$} = -\ln E_t \left[M_{t,t+1} \frac{\Pi_t}{\Pi_{t+1}} \right]. \quad (32)$$

The upper half of Figure 4 shows the calibrated model's nominal risk-free rate and price-dividend ratio as functions of the real state-variables, $\sigma_{l,t}^2$ and $\sigma_{p,t}^2$. Two standard deviation changes in the volatility of the relative price growth is given on the x-axis, while variation in the conditional volatility of luxury good consumption growth is shown by changing the

¹⁷When analyzing relevant time-series moments of the price-dividend ratio in the model, we construct this variable to be comparable with how it is constructed in the data (see, e.g., Boudoukh, Michaely, Robertson, and Richards, 2007). I.e., the quarterly dividends are reinvested until the end of the year to create annual dividends, as is the case for annual dividends imputed from the CRSP ex- and cum-dividend market returns. We do this every quarter to arrive at an annual price-dividend ratio series overlapping at the quarterly frequency.

plotted line from solid to dashed. In particular, the dashed line corresponds to times when luxury growth volatility is two standard deviations above average, while the solid line corresponds to when it is two standard deviations below its average value. Since the conditional volatility of luxury good consumption growth is a counter-cyclical business cycle frequency variable, the dashed line is labeled "recession" while the solid line is labeled "expansion".

The risk-free rate is increasing in the conditional volatilities, indicating that bonds are risky. I.e., bonds fall in value when there are negative shocks to the risk factors. This is similar to the model in Wachter (2006), where the risk-free rate is positively related to the conditional volatility of surplus consumption, which in turn is negatively correlated with aggregate consumption shocks. The nominal risk-free rate is of course increasing in inflation (not shown in graph), but since inflation is close to uncorrelated with the risk factors shocks to inflation are not important for nominal bond risk premia in the model. The price-dividend ratio is decreasing in both volatilities. Dividend growth is unpredictable in the calibrated model, so movements in the price-dividend ratio are solely due to movements in discount rates (Campbell and Shiller, 1988). Both the equity risk premium and the risk-free rate contribute to increasing discount rates when the conditional volatilities of the risk factors are high. The lower half of Figure 4, shows that both the equity risk premium and the risk-premium on a 5-year risk-free zero-coupon bond are increasing in the real state variables. Since the two state-variables operate at difference frequencies, asset risk premia in general inherit the two-frequency dynamics. Note from the plot, however, that the fluctuations in the equity risk premium are larger at the generational cycles followed by $\sigma_{\Delta p}^2$ than over the business cycle, as measured by fluctuations in $\sigma_{\Delta l}^2$. This previews the result that return predictability, in the model and in the data, is stronger at frequencies lower than the business cycle frequency than at the business cycle frequency.

Figure 5 shows how the conditional volatility and Sharpe ratio of equity returns vary with the state variables. Note that equity return volatility is only slightly higher in recessions (dashed line) than in expansions (solid line), but instead varies more with the lower frequency variation in risk. This is in part because the volatility of dividends is positively related to this slower moving state variable, and in part because changes in the conditional volatility of more persistent shocks have a bigger impact on the volatility of the equity price. The equity Sharpe ratio, however, is much more closely related to the business cycle; high in recessions and low in expansions. Time-variation in the Sharpe ratio is driven by time-variation in the conditional volatility of the stochastic discount factor. The fluctuations in the conditional volatility of luxury good consumption are estimated to be larger in absolute

Figure 4

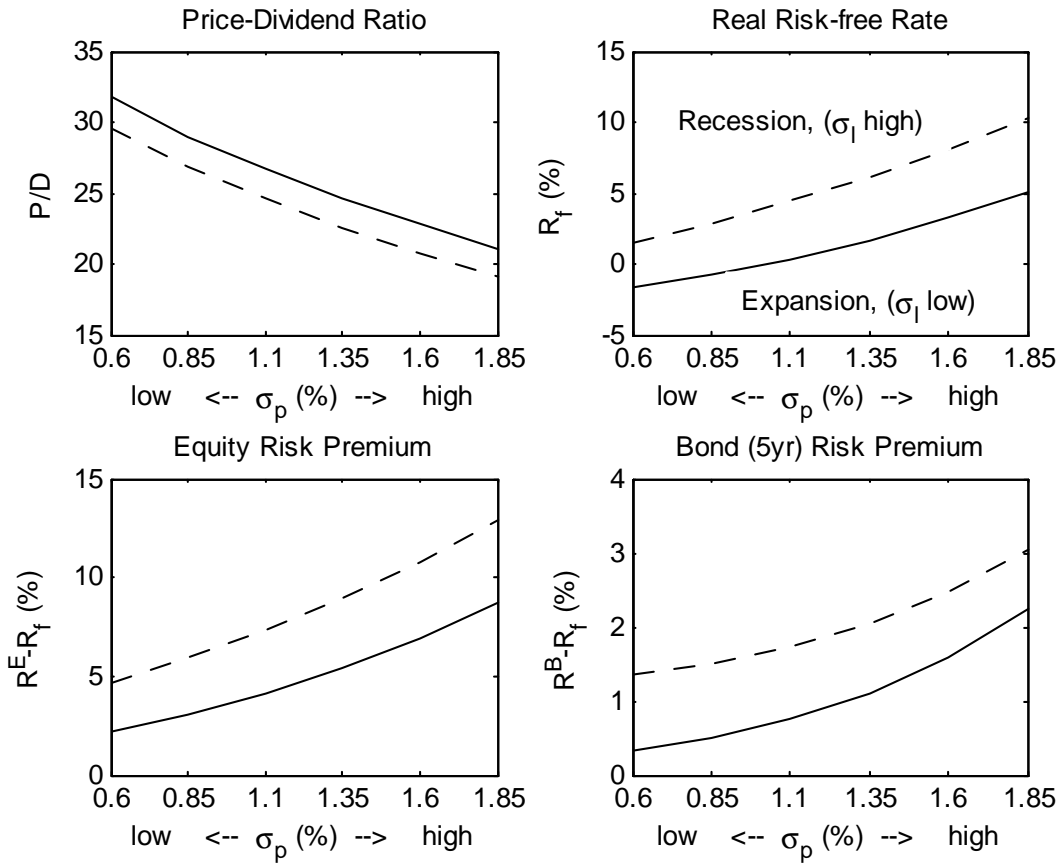


Figure 4: The figure gives two standard deviation fluctuations in the conditional variance of the relative price, $\sigma_{p,t}^2$, on the x-axis, and two standard deviations in the conditional variance of luxury good consumption growth, $\sigma_{l,t}^2$, as a dashed line (high) and a solid line (low). The plots are of the equity market price-dividend ratio, the real risk-free rate, the equity premium, and the risk premium on a 5-year nominal zero-coupon bond. All variables are solved for at a quarterly frequency and annualized.

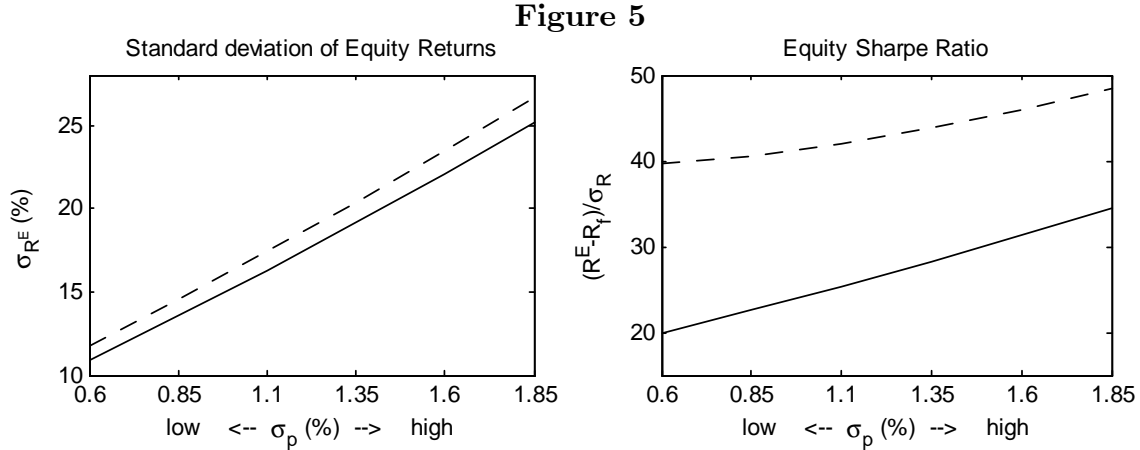


Figure 5: The figure gives two standard deviation fluctuations in the conditional variance of the relative price, $\sigma_{p,t}^2$, on the x-axis, and two standard deviations in the conditional variance of luxury good consumption growth, $\sigma_{l,t}^2$, as a dashed line (high) and a solid line (low). The plots are of the volatility of equity market returns and the equity Sharpe ratio. All variables are solved for at a quarterly frequency and annualized.

value than the fluctuations in the conditional volatility of relative price growth. Since the former operates at a business cycle frequency, the Sharpe ratio is more tightly related to business cycle fluctuations than the lower frequency, generational fluctuations. However, due to the volatility effect, the equity risk premium has a stronger relation to these low frequency movements in risk. This difference in persistence between the equity Sharpe ratio and the equity risk premium in the model is consistent with the empirical findings in Lettau and Ludvigson (2007).

5.2 Unconditional Moments

Panel A of Table 4 gives a summary of the unconditional simulated sample moments of the model's exogenous state processes versus the data, while Panel B shows the unconditional asset pricing moments in the data as well as in simulated data from the model. The simulated moments are computed as sample averages of 1,000 samples of the same length as in the data (191 quarters). Thus, any small sample issue in the moments from the data is replicated in the model generated moments. To evaluate whether the model is able to match the moments in the data, the moments from the model are given as the median, the 5th percentile and 95th percentile moments from the 1,000 samples. Thus, the table effectively reports the 90% confidence interval of each moment.

Table 4 - Unconditional Moments (Annual)

Table 4: This table reports the median, 5th percentile and 95th percentile moments from small-sample model simulations, as well as the corresponding moment value in the data. Panel A reports relevant moments for the state processes in the model. Panel B reports asset pricing moments. The sample period is 1959Q3 - 2006Q4; 190 quarterly observations.

Panel A: Summary of Unconditional State Processes Moments

	<i>Data</i>	<i>Model</i>		<i>Data</i>	<i>Model</i>
$E(\Delta l)$ (%)	2.37	1.24	$E(\Delta d)$	2.58	2.58
$\sigma(\Delta l)$ (%)	4.89	4.89	$\sigma(\Delta d)$	12.07	12.07
$E(\Delta p)$ (%)	0.27	0.27	$E(\Delta \pi)$	3.81	3.81
$\sigma(\Delta p)$ (%)	2.32	2.32	$\sigma(\Delta \pi)$	1.31	1.31

Panel B: Financial Moments

	<i>Data</i>	<i>Model simulated samples</i>		
		<i>Median</i>	<i>5th-percentile</i>	<i>95th-percentile</i>
$E(R^E - R_f)$ (%)	5.93	6.36	2.64	9.30
$\sigma(R^E - R_f)$ (%)	16.83	17.33	14.22	22.98
$SR(R^E - R_f)$	0.35	0.36	0.13	0.58
$E(P/D)$	27.28	25.95	23.74	28.18
$\sigma(p-d)$	0.27	0.10	0.06	0.18
$E(R_f^{\text{real}})$ (%)	2.12	1.97	0.87	3.91
$\sigma(R_f^{\text{real}})$ (%)	1.35	1.11	0.70	2.39
$E(R_f)$ (%)	5.93	5.74	3.98	8.07
$\sigma(R_f)$ (%)	1.52	1.56	1.10	2.53
$E(R_{3yr} - R_f)$ (%)	0.76	1.48	0.31	2.41
$\sigma(R_{3yr} - R_f)$ (%)	3.55	5.44	4.41	7.62
$SR(R_{3yr} - R_f)$	0.21	0.27	0.05	0.50
$E(R_{5yr} - R_f)$ (%)	1.10	1.75	0.12	3.03
$\sigma(R_{5yr} - R_f)$ (%)	5.77	7.31	5.99	9.98
$SR(R_{5yr} - R_f)$	0.19	0.24	0.01	0.46
$E(R_{10yr} - R_f)$ (%)	1.10	1.89	-0.16	3.52
$\sigma(R_{10yr} - R_f)$ (%)	8.13	9.23	7.69	12.08
$SR(R_{10yr} - R_f)$	0.14	0.21	-0.01	0.42

The median sample equity premium and return volatility are close to the historical values in the data. Thus, the model is able to account for the equity premium puzzle (Mehra and Prescott, 1985). This is consistent with the empirical findings of Ait-Sahalia, Parker, and Yogo (2004), who estimate a low level of relative risk aversion using a different measure of luxury good consumption as discussed earlier. Note that the volatility of equity returns is greater than that of dividend growth. This "excess" return volatility (Shiller, 1981) is due to time-variation in discount rates. The model matches relatively well the level of the equity market price-dividend ratio, but its volatility is too low compared to the data. This indicates that there is more time-variation in discount rates, expected dividend growth, or both, in the data than what is generated by the model. We will discuss this further when evaluating the small-sample excess return predictability generated by the model.

The model also matches well the unconditional level and volatility of the nominal and the real risk-free rates. Excess return and return volatility to long-term bonds are increasing in maturity, as in the data, and again the model matches well the levels of these variables, with the exception of the volatility of 3- and 5-year zero-coupon bond returns, which is high relative to the data. However, the 10-year bond return volatility match the data well. The Sharpe ratios of bonds in the model are less than the Sharpe ratio of equities and decreasing in bond maturity. This pattern is consistent with the historical data.

In sum, the model matches the first and second moments of aggregate equity returns and nominal risk-free bonds quite well. This is especially noteworthy considering that the dynamic behavior of the habit level is not calibrated to match these asset pricing moments, but instead estimated to be consistent with the observable macro economic data. Before we turn to the dynamic behavior and conditional moments of the model, however, we investigate whether the equity premium implied by the model calibration is consistent with an alternative empirical measure of the model implied equity premium.

5.2.1 Model Implied Covariances and the Data

The model calibration with respect to the equity risk premium relies on an estimate of the correlation between real dividend growth and the two real risk factors in the economy. Dividends are, however, difficult to measure accurately (see, e.g., Boudoukh et al, 2007). We therefore in this section consider a different measure of the risk premium: the covariance of excess equity returns with the stochastic discount factor. Given that we have an *observable* stochastic discount factor, we can investigate its joint dynamic properties with excess equity

market returns directly.

The unconditional Law of One Price (see, e.g., Campbell, 1999) states that

$$E [R_{i,t+1}^e] = -R_f Cov (M_{t+1}, R_{t+1}^e) \quad (33)$$

$$\Downarrow$$

$$\frac{E [R_{i,t+1}^e]}{\sigma (R_{t+1}^e)} = -R_f \rho (M_{t+1}, R_{t+1}^e) \sigma (M_{t+1}), \quad (34)$$

where $\rho(\cdot)$ denotes correlation and $\sigma(\cdot)$ denotes standard deviation. Given that the risk-free rate is small, a high unconditional risk premium must come from a large negative unconditional covariance between excess returns and the stochastic discount factor (equation (33)). A large Sharpe ratio must, per equation (34), come from a negative correlation between excess returns and the stochastic discount factor, as well as a high volatility of the stochastic discount factor.

The annualized volatility of the stochastic discount factor in the data, given the calibrated values of γ and δ and the time-series properties of Δl and Δp , is 51%. The average real risk-free rate is 2%, and so the maximum Sharpe ratio given by the model is 0.52 (annual).¹⁸ The quarterly correlation of contemporaneously reported consumption, price and returns data, gives a correlation between excess returns and the stochastic discount factor of only -0.14 . This implies an equity Sharpe ratio of only 0.07. Using, instead, end-of-year data consumption data from 1959Q4 - 2006Q4 and annual returns data 1959Q3 - 2006Q3 (i.e., annual data using the beginning of quarter consumption timing convention of Campbell, 1999, to account for time-aggregation), the correlation between the stochastic discount factor and excess equity market returns is -0.48 , which implies an equity Sharpe ratio of 0.25. This is closer to the 0.35 found in the data and implied by the model calibration.

More fundamentally, these differences in measured correlation may indicate the presence of measurement error (e.g., Christiano, Eichenbaum, and Marshall, 1991; Wilcox, 1992; Working, 1960). Measurement error in the consumption and price data may lead to economically significant understatement of the true correlation between the real risk factors and asset returns. Further, seasonality in consumption spending (Jagannathan and Wang, 2007) or frictions to adjusting consumption behavior (e.g., Lynch, 1996; Gabaix and Laibson, 2002) are not explicitly accounted for in the simple model in this paper, but may be present in the data. However, Parker and Julliard (2005) suggest a way to test consumption-based models

¹⁸The maximal Sharpe ratio is given by $\sigma(M)/E(M)$ (see, e.g., Hansen and Jagannathan, 1991).

which is robust to some of these frictions and data problems. We adopt their methodology as an alternative way of investigating whether the model is able to empirically address the equity premium puzzle.

Using equations (6) and (7), we can write the stochastic discount factor in terms of the marginal utility of the representative agent:

$$M_{t+1} = \beta \frac{MU_{t+1}}{MU_t}. \quad (35)$$

Since, by no-arbitrage, the S period risk-free rate at time $t + 1$ is $R_{f,t+1}^S = 1/E_{t+1}[M_{t+1+S}]$, we have that $MU_{t+1} = E_{t+1}[\beta^S MU_{t+1+S} R_{f,t+1}^S]$. Next, define

$$M_{t+1}^S \equiv \beta^{1+S} R_{f,t+1}^S \frac{MU_{t+1+S}}{MU_t} = \beta^{1+S} R_{f,t+1}^S \left(\frac{L_{t+1+S}}{L_t} \right)^{-\gamma} \left(\frac{P_{t+1+S}}{P_t} \right)^{-\delta}. \quad (36)$$

Using the Law of Iterated Expectations, we then have

$$E[R_{t+1}^e] = -Cov(M_{t+1}^S, R_{t+1}^e) / E[M_{t+1}^S]. \quad (37)$$

This version of the stochastic discount factor holds for all horizons S , and thus allows us to look at the correlation between longer horizon growth rates in the relative price and luxury good consumption versus 1 quarter equity returns. Parker and Julliard (2005), shows that doing so can lead to quite different quantitative inferences from the model. Table 5 shows the sample measure of the risk premium, constructed using equation (37) and the data sample in this paper for S ranging from 0 (no timing issues) to 16 quarters. The focus here is on the point estimates of the model's equity risk premium, and we therefore do not report standard errors of the estimated values. The equity risk premium calculated using equation (37) above starts at 0.95% for $S = 0$, and peaks at 6.63% for $S = 4$. Thus, the sample equity risk premium of 5.93% is well within the maximum the model empirically is measured to deliver. For comparison, the same measure of the risk premium but using the standard Consumption CAPM (i.e., $M_{t+1}^S = \beta^S (C_{t+1+S}/C_t)^{-\gamma}$) is reported in the rightmost column. We assume $\gamma = 10$ for this benchmark model as well. In this case, the model estimated risk premium starts 0.30% at $S = 0$ and peaks at 1.41% at $S = 8$. Thus, one would need $\gamma = 42$ for the standard model to match the historical average excess equity return at the "optimal" horizon, S . However, this high value of relative risk aversion would lead to a risk-free rate puzzle (Weil, 1989).

Table 5 - Risk Premium Estimate from Empirical Covariance with SDF

Table 5: This table reports estimates of the equity risk premium using both the model in this paper and the standard Consumption CAPM. The estimated equity premium uses the relation $E(R_{t+1}^e) = -Cov(M_{t+1}^S, R_{t+1}^e)/E(M_{t+1}^S)$, which comes from the law of one price. The data is quarterly and reported risk premiums are annualized by simply multiplying the quarterly risk premium by 4. The sample period is 1959Q3 - 2006Q4.

	Two-factor model $M_t^S = \beta^{S+1} e^{-\gamma \Delta l_{t+S} - \delta \Delta p_{t+S}}$	Standard CCAPM $M_t^S = \beta^{S+1} e^{-\gamma \Delta c_{t+S}}$
$\hat{E}(R_{mkt}^e S = 0)$ (%)	0.95	0.30
$\hat{E}(R_{mkt}^e S = 1)$ (%)	2.46	0.76
$\hat{E}(R_{mkt}^e S = 4)$ (%)	6.63	1.10
$\hat{E}(R_{mkt}^e S = 8)$ (%)	5.28	1.41
$\hat{E}(R_{mkt}^e S = 12)$ (%)	2.36	0.93
$\hat{E}(R_{mkt}^e S = 16)$ (%)	0.79	0.37

Thus, judging from returns data directly, and not relying on poorly measured aggregate cash flow data, the model is able to empirically match the equity risk premium. However, this success relies on allowing for longer horizon growth rates for the real factors in the stochastic discount factor, as in Parker and Julliard (2005). Taking the model as is, this empirical difference in performance over different time-horizons must be due to measurement error in the consumption and relative price data.

5.3 Conditional Moments

In this section, we will consider the conditional properties of the model. In particular, since the model has two state variables that operate at different frequencies, we investigate whether this helps the model get the relative persistence of key economic state variables, such as the risk-free rate, the price-dividend ratio, and the term spread right. Further, we follow the standard in the literature and look at forecasting regressions of future excess stock and bond returns (e.g., Fama and French, 1989) using both historical data and data simulated from the calibrated model.

Table 4 - Annual Autocorrelations

Table 6: This table reports the median, 5th percentile and 95th percentile moments from small-sample model simulations of annual autocorrelations of key variables, as well as the corresponding moment value in the data. The sample period is 1959Q3 - 2006Q4. † The annual autocorrelation of the conditional equity Sharpe ratio is taken from Lettau and Ludvigson (2007).

<i>Annual autocorrelation of variable:</i>	<i>Data</i>	<i>Model simulated samples</i>		
		<i>Median</i>	<i>5th-percentile</i>	<i>95th-percentile</i>
$\sigma_t(\Delta l_{t+1})$	0.32	0.31	0.13	0.51
$\sigma_t(\Delta p_{t+1})$	0.73	0.65	0.46	0.82
$SR_t(R_{t+1}^E - R_{f,t})^\dagger$	0.52	0.46	0.28	0.64
$p - d$	0.87	0.73	0.57	0.88
R_f^{real}	0.55	0.51	0.31	0.71
R_f	0.76	0.60	0.39	0.76
$y_{3yr} - y_{1yr}$	0.46	0.52	0.29	0.70
$y_{5yr} - y_{1yr}$	0.51	0.55	0.32	0.72
$y_{10yr} - y_{1yr}$	0.49	0.58	0.36	0.74

5.3.1 Persistence

The Sharpe Ratio. Time-variation in the conditional Sharpe ratio of the an asset must be due to time-variation in the correlation between the asset’s returns and the stochastic discount factor, time-variation in the risk-free rate, or time-variation in the conditional volatility of the stochastic discount factor.

$$\frac{E_t(R_{i,t+1} - R_{f,t})}{\sigma_t(R_{i,t+1})} = -\rho_t(R_{i,t+1}, M_{t+1}) \sigma_t(M_{t+1}) R_{f,t}. \quad (38)$$

In the model, the latter is the important channel for the conditional Sharpe ratio of the aggregate equity market. Table 6 shows that the median annual autocorrelation of the conditional Sharpe ratio of equities is 0.46, which is close to the 0.52 as measured empirically by Lettau and Ludvigson (2007). The two real state variables in the model, $\sigma_{\Delta l_t}$ and $\sigma_{\Delta p_t}$ have annual autocorrelations of 0.32 and 0.76, respectively, and the dynamic behavior of the Sharpe ratio is in the middle of these as both are important for the conditional volatility of the stochastic discount factor. Further, as shown in Figure 5, the Sharpe ratio is strongly

counter-cyclical.

The Price-Dividend Ratio. The price-dividend ratio is more persistent than the Sharpe ratio, both in the model and in the data. The historical annual autocorrelation of the aggregate log price-dividend ratio is 0.87. The model can, however, only just match this high persistence. In particular, the 95th percentile annual autocorrelation of the dividend yield in the model is 0.883, while in the data it is 0.868.¹⁹ While this persistence is within the model's 95th percentile, we would like the model to more comfortably match this important moment. In fact, there is reason to think this number is understated in the model. In particular, the autocorrelation of the price-dividend ratio in the model is driven mainly by the persistence of the most persistent state variable, the conditional volatility of relative price growth. This variable's quarterly autocorrelation is 0.934, which is 0.761 annually. However, the estimate of the persistence parameter $\beta_{p,3}$ in the maximum likelihood procedure applied in section 4 is downward biased (see, e.g., Bollerslev and Wooldridge, 1992). It is well-known that estimates of persistence is downward biased in small samples. This is also clear from Table 6. The estimated conditional volatility of relative price growth has an autocorrelation of 0.73. However, when using the estimated value of $\beta_{p,3}$ in the model calibration, the model generates a median annual autocorrelation of the conditional volatility of the relative price of only 0.65, due to the small-sample bias. We have not made any attempt to correct for small-sample bias in the estimation of the parameters of the state processes, and thus the downward bias in the estimated persistence of the state variables is inherited by the model's dividend yield. Thus, the difference between the persistence of the dividend yield in the data and in the model reported in Table 6 is biased upwards.

The fact that the Sharpe ratio and the dividend-price ratio in the model, as in the data, operate at quite different frequencies, is a departure from current benchmark models where only one state variable drives the dynamic behavior of risk (e.g., Campbell and Cochrane, 1999; Wachter, 2006).

The Risk-Free Rate. The model matches well the persistence of the real risk-free rate and the slope of the term structure, measured as the difference between the 3-, 5-, and 10-year yield and the 1-year yield on nominal risk-free zero-coupon bonds. The term spreads have a significant business cycle component, as in the data. Importantly, the persistence of

¹⁹Both of these autocorrelations, however, are downward biased given the relatively small sample. The population autocorrelation of the dividend yield in the model is 0.81.

the yield spreads are significantly lower than the persistence of the dividend yield, which indicates that the former are more tightly linked to business cycle fluctuations, while the latter is more strongly related to generational fluctuations in discount rates. Again, this is a feature of the data that cannot be replicated in a model with only one risk factor.

The nominal risk-free rate, however, is only barely as persistent as in the data. Since the persistence of the real risk-free rate is matched in the model, this indicates that the assumed inflation process in the model is not rich enough to match the persistence of the nominal short-rate. We leave a more in-depth analysis of the inflation dynamics to future research.

In sum, the model generates economic state variables with different persistence, consistent with the historical data. This differences in persistence comes from the *estimated* difference in persistence of the conditional volatilities of the two real risk factors in the model, which operate at the business cycle frequency and an even lower "generational" frequency.

5.3.2 Forecasting Regressions

Forecasting regressions of excess equity and bond returns are a standard way to detect time-variation in equity and bond risk premia (see, e.g., Fama and French, 1989, Campbell and Shiller, 1988). We consider forecasting regressions of excess equity and bond returns using simulated data from the model and compare this to forecasting regressions using historical data. The forecasting regressions are standard. Let $r_{i,t,t+q}^e$ be the log return on asset i in excess of the risk-free rate from time t to $t + q$ and let x_t be a vector of forecasting variables observable at time t . The generic forecasting regressions we run are of the form:

$$r_{i,t,t+q}^e = \alpha_i + \beta_i' x_t + \varepsilon_{i,t,t+q}. \quad (39)$$

When reporting regressions using simulated data from the model, the t -statistics, the regression coefficients, and the R^2 's are sample average of 1,000 regressions of the same size as those in the data (191 quarters). This way, the small-sample issues that often arise in these regressions (e.g., Stambaugh, 1999; Campbell and Yogo, 2003) are replicated in the model regressions. All t -statistics are corrected for heteroskedasticity and autocorrelation using Newey-West standard errors, in both historical and simulated data regressions.

Equity returns. Risk premia are in the model driven by the conditional volatility of luxury good consumption growth and relative price growth. In particular, equation (14) shows that the maximal conditional price of risk in the economy is a function of $\sigma_{l,t}^2$, $\sigma_{p,t}^2$,

and the interaction term $\sigma_{l,t}\sigma_{p,t}$. Panel A of Table 7 shows how excess equity returns at increasing forecasting horizons are related to these variables in the model in a univariate setting. In particular, while $\sigma_{l,t}^2$ is positively related to the risk premium, per Figure 4, the model regressions show that in samples of the same size as that we have available, this business cycle predictability is not statistically significant when using $\sigma_{l,t}^2$ alone. The R^2 is increasing in the horizon, but small at only 1.2% at the annual forecasting horizon. Note, however, that an R^2 of 1.2%, implies a risk premium standard deviation of 1.9% given the model's realized equity return volatility. Thus, as noted by Kandel and Stambaugh (1996), small R^2 's in these regressions does not imply that the time-variation in the risk premium is *economically* insignificant. However, the R^2 's in the data are effectively zero. Since $\sigma_{l,t}^2$ is estimated there is an errors-in-variables problem that biases the regression coefficients and R^2 's downwards.

Both in the model and in the data, however, $\sigma_{p,t}^2$ is a significant predictor of future excess equity returns. The regression coefficient and R^2 's are increasing in the return forecasting horizon. The fact that $\sigma_{p,t}^2$ is a stronger predictive variable than $\sigma_{l,t}^2$ is expected given Figure 4, which shows that the equity risk premium fluctuates more at the generational frequency of $\sigma_{p,t}^2$.

Campbell and Shiller (1988) note the special role of the dividend-price ratio in predicting future excess returns, and Fama and French (1989) emphasize that the dividend price ratio is a lower frequency predictive variable compared to the business cycle related term spread. The rightmost columns of Panel A in Table 7 shows that the dividend yield does indeed forecast future excess equity returns both in the data and in the model. However, the forecasting power of the standard measure of the dividend yield employed here (cash-dividends imputed from the CRSP files, which quarterly data is readily available for) is quite weak. This is consistent with the results reported in Boudoukh et al. (2007). As expected, the regression coefficient and the R^2 are increasing in the return horizon also here.

Importantly, the dividend yield is more slow-moving than the business cycle, which indicates there may be business cycle dynamics in the risk premium not picked up by the dividend yield. Panel B of Table 7 shows the multivariate regression using $\sigma_{l,t}^2$, dp_t , and the interaction term $\sigma_{l,t}^2 dp_t$, where dp_t is a measure of the dividend yield. We use both the cash dividend yield and the *CAY*-variable of Lettau and Ludvigson (2001) as measures of the aggregate equity market dividend yield. The latter variable is a measure of the aggregate consumption-wealth ratio that Lettau and Ludvigson (2001) show is a strong predictive variable. The model predicts that the risk premium is increasing in recessions, as measured by

$\sigma_{i,t}^2$, but more so when the dividend yield is otherwise high. This latter effect is reflected in a positive coefficient on the interaction term. In simulated data from the model, the inclusion of the business cycle variable $\sigma_{i,t}^2$ increases the adjusted R^2 of the dividend yield regressions from 2.6% to 6.9% in the annual regressions, but on average this effect is not statistically significant in the model regressions. In the data, however, the interaction terms come in significant for both the cash–dividend yield and the *CAY*-variable, if only at the quarterly horizon for the latter. The effect of the business cycle variable decreases for forecasting horizons longer than a year, as could be expected since the business cycle has a relatively short half-life. Overall, the model delivers on average somewhat too little predictability relative to that found in the data. This is consistent with the fact that the volatility of the price–dividend ratio in the model is smaller than that in the data.

Next, we investigate further the existence of a business cycle component in expected excess equity market returns. In particular, Panel C of Table 7 shows the dividend yield and *cay* forecasting regressions interacted with an *NBER* recession indicator and lagged *GDP* growth. The former is not strictly a valid forecasting variable, as the *NBER* dates are set ex-post, but it is still interesting in order to pin this effect to the business cycle. Again, recessions (i.e., below average lagged *GDP* growth, or a positive *NBER* indicator; note that the signs will thus be different) are associated with a high risk premium, and the effect a recession has on the risk premium is increasing in the level of the dividend yield.

Predictability of aggregate, excess equity returns are well documented (e.g., Fama and French, 1988; Campbell and Shiller, 1988; Lettau and Ludvigson, 2001). Fama and French (1989), in particular, emphasized that the dividend price ratio is a lower frequency predictive variable compared to the business cycle related term spread. The results reported here can thus be seen as giving a macro economic, consumption-based foundation for the existence of risk factors operating at both of these frequencies.

The strong predictive ability of the conditional volatility of the relative price is especially noteworthy given that this variable does not have a market price component, as opposed to what is the case for both the dividend yield and the *CAY*-variable. Figure 6 shows the historical dividend yield and *CAY*-variable versus the estimated conditional volatility of the relative price growth. Both variables are, consistent with the model, empirically significantly and positively correlated with the conditional volatility of relative price growth. *NBER* style recession indicators (vertical bars) are plotted and indirectly show that the predictive power of this variable is not mainly due to business cycle dynamics, but to a lower frequency

Forecasting Excess Equity Market Returns

Table 7: Panel A of this table reports average forecasting regression coefficients when regressing the quarterly, annual and 4-year excess equity returns on the lagged luxury good consumption and relative price growth volatilities, $\sigma_{I,t}^2$ and $\sigma_{p,t}^2$, as well as the log annual dividend-yield (dpt_t). The numbers from the model are based on 1000 simulated quarterly samples of 191 observations. Panel B reports results from multivariate regression using the conditional variance of luxury good consumption growth as a business cycle interaction variable. All individual regressors have been demeaned in the interaction terms. Panel C reports similar regressions as in Panel B, but with other, standard business cycle variables. $gdpt_t$ denotes quarterly, real, per-capita GDP growth, while $nber_t$ denotes an NBER recession indicator. The cay_t variable is the measure of the aggregate consumption-wealth ratio from Lettau and Ludvigson (2001). Heteroskedasticity and autocorrelation corrected t -statistics are given in parenthesis. The sample period in panels A and B is 1959Q3 - 2006Q4, while in Panel C the sample period is 1952Q1 - 2006Q4.

Panel A: Forecasting Regressions: Lux. Volatility, Price Volatility, and the Dividend Price ratio as univariate predictive variables

$Horizon$ (q)	$\sigma_{I,t}^2 \times 100$		$\sigma_{p,t}^2 \times 100$		dpt_t	
	Model β	Data R_{adj}^2	Model β	Data R_{adj}^2	Model β	Data R_{adj}^2
1	0.120 (0.524)	0.1% (0.502)	0.022 (0.502)	0.0% (0.502)	1.346 (1.280)	1.2% (1.487)
4	0.326 (0.545)	1.2% (0.054)	-0.059 (0.054)	0.0% (0.054)	4.567* (1.807)	5.0% (1.962)
16	0.524 (0.439)	2.7% (0.243)	0.429 (0.243)	0.0% (0.243)	11.416** (2.206)	11.7% (2.771)
					0.098 (1.332)	1.2% (1.144)
					0.334* (1.795)	5.0% (1.360)
					0.553** (2.139)	12.3% (0.521)
					0.069 (1.360)	2.6% (1.360)
					0.149** (0.521)	5.2% (0.521)

Panel B: Forecasting Regressions: The Dividend Yield, cay and the Business Cycle

$Horizon$ (q)	Model		Historical data (1959Q3 - 2006Q4; 190 observations)		cay_t and $\sigma_{I,t}^2$	
	dpt_t	$\sigma_{I,t}^2 dpt_t$	dpt_t	$\sigma_{I,t}^2 dpt_t$	cay_t	$\sigma_{I,t}^2 cay_t$
1	0.056 (0.243)	0.102 (1.347)	0.474 (0.227)	1.125 (1.353)	1.691** (4.270)	67.567** (2.333)
4	0.112 (0.187)	0.339* (1.800)	0.786 (0.209)	3.936** (2.343)	-0.478 (-0.453)	55.220 (1.631)
					0.096 (0.241)	6.0% (2.333)
					0.096 (0.241)	6.0% (2.333)
					5.001** (3.622)	15.6% (1.631)

Panel C: Forecasting Regressions: The Dividend Yield, cay and other Business Cycle variables in *longer* sample

$Horizon$ (q)	Historical data (1952Q1 - 2006Q4; 220 observations)		Historical data (1959Q3 - 2006Q4; 190 observations)		$gdpt_t$ and $\sigma_{I,t}^2$	
	dpt_t	$nber_{t-1} dpt_t$	dpt_t	$gdpt_t dpt_t$	$gdpt_t$	$\sigma_{I,t}^2 gdpt_t$
1	0.002 (0.101)	0.020 (1.494)	0.019 (1.472)	-2.884* (1.677)	1.418** (4.108)	-96.507** (-2.081)
4	0.003 (0.103)	0.085** (1.984)	0.077* (1.785)	-9.338** (-2.511)	4.852** (3.637)	-86.619 (-0.7844)
					1.3% (1.677)	6.2% (-2.081)
					10.4% (-1.944)	18.4% (-0.7844)

Figure 6

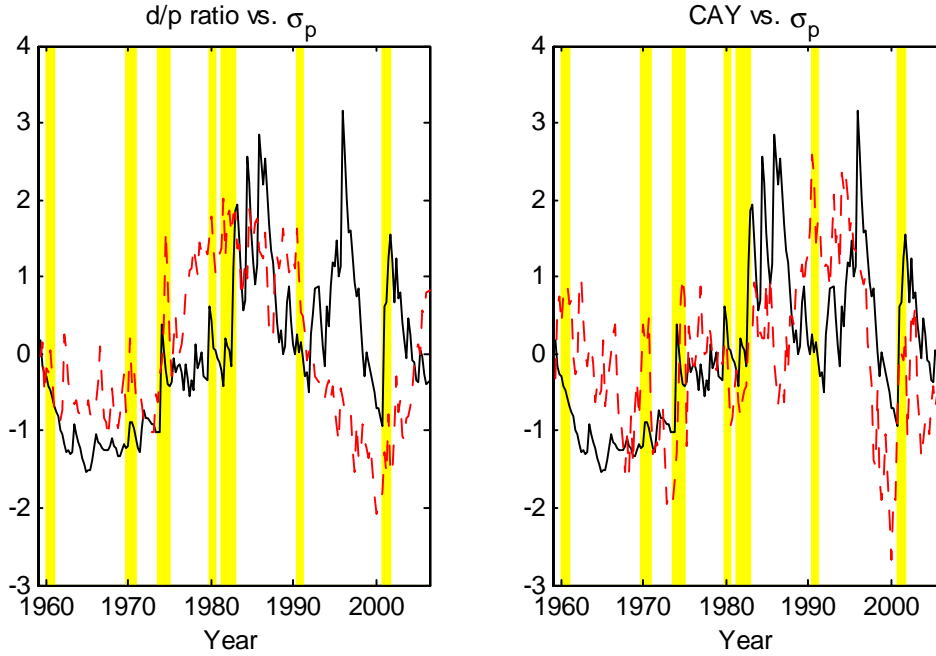


Figure 6: The figure shows the estimated historical conditional volatility of relative price growth, $\sigma_{p,t}$, (solid line) versus the aggregate historical log dividend yield and the *cay*-variable of Lettau and Ludvigson (2001). The bars are NBER recession indicators. The sample is 1959Q3 - 2006Q4.

component.

Bond returns. Table 8 shows forecasting regressions with excess bond returns as the dependent variable. The bonds are nominal, default-free zero-coupon bonds with maturities 2, 3, 4, and 5 years obtained from the CRSP files. In all cases we consider an annual return forecasting horizon, where the data is overlapping at a quarterly frequency. Again, the reported regression coefficients are sample averages of 1,000 regressions of the same length as the data sample, but using simulated data from the model. The reported regression *t*-statistics are the average Newey-West *t*-statistics corrected for the length of the observation overlap.

Table 8 shows that the luxury good volatility is not a good predictor of future excess bond returns in the model or in the data. The regression coefficient is increasing with maturity for maturities up to 4 years, but then decreases both in the data and in the model. Business cycle fluctuations are relatively fast-moving and therefore less important for longer maturity bonds. The volatility of the relative price, however, is the more successful forecasting variable. It

TABLE 8
Forecasting Annual Excess Bond Returns

Table 8: Panel A of this table reports average forecasting regression coefficients when regressing annual excess bond returns on 2-, 3-, 4-, and 5-year bonds on the lagged luxury good consumption and relative price growth volatilities, $\sigma_{l,t}^2$ and $\sigma_{p,t}^2$, as well as the Cochrane-Piazzesi factor. The numbers from the model are based on 1000 simulated quarterly samples of 191 observations. Heteroskedasticity and autocorrelation corrected t -statistics are given in parenthesis. The sample period is 1959Q3 - 2006Q4.

<i>Maturity</i> (<i>yrs</i>)	$\sigma_{l,t}^2 \times 100$				$\sigma_{p,t}^2 \times 100$				C&P factor			
	Model		Data		Model		Data		Model		Data	
	β	R_{adj}^2	β	R_{adj}^2	β	R_{adj}^2	β	R_{adj}^2	β	R_{adj}^2	β	R_{adj}^2
2	0.042 (0.857)	0.3%	0.046 (0.341)	0.0%	0.302* (1.661)	1.0%	0.542** (2.152)	4.8%	0.712** (3.011)	4.2%	0.446** (7.690)	30.4%
3	0.046 (0.691)	0.2%	0.0784 (0.286)	0.0%	0.406 (1.621)	1.0%	0.849* (1.749)	3.3%	0.948** (2.976)	4.0%	0.847** (8.179)	32.2%
4	0.046 (0.585)	0.1%	0.108 (0.275)	0.0%	0.480 (1.586)	0.9%	1.208* (1.726)	3.5%	1.112** (2.942)	3.8%	1.208** (8.455)	34.3%
5	0.045 (0.514)	0.1%	0.050 (0.100)	0.0%	0.532 (1.557)	0.9%	1.498* (1.715)	3.6%	1.229** (2.912)	3.7%	1.498** (8.763)	34.8%

forecasts excess bond returns in the data with a significant coefficient at all horizons. The regression coefficient is increasing with bond maturity as the bonds with longer maturity are more affected by this very persistent risk factor. The R^2 's are slightly decreasing in the maturity of the bonds. The regressions using model simulated data are consistent with these findings, except that the average forecasting power of $\sigma_{p,t}^2$ is not significant for bond maturities of more than 2 years.

In the model, the R^2 's of the bond return predictability are fairly low. An R^2 of 0.9% for the 5-year bond, however, implies that the standard deviation of expected excess 5-year bond returns in the model is 0.6% p.a, which is still economically significant relative to the average risk premium of 1.1% p.a. However, the empirical regressions find higher R^2 's around 4%, which implies an annual standard deviation of the 5-year bond risk premium of 1.1%. We should note that there is no significant small-sample bias in the regressions using the estimated volatilities as the forecasting variables. Thus, again there appears to be somewhat stronger predictability in the data than that delivered by the model on average.

Finally, we consider the Cochrane and Piazzesi (2005) factor as a forecasting variable. This is the currently most successful forecasting variable of bond returns and it delivers powerful empirical evidence of the failure of the expectations hypothesis. Of course, we already know that the expectations hypothesis fails in the model, as bond risk premia are time-varying (see Fama and Bliss, 1986; Campbell and Shiller, 1988), but using this factor provides a benchmark for assessing the magnitude of this failure in the model relative to the data. Cochrane and Piazzesi (2005) run the following regression

$$\overline{r\bar{x}}_{t+1} = \gamma' f_t + \varepsilon_{t+1}, \quad (40)$$

where $\overline{r\bar{x}}_{t+1}$ is the average excess annual return on 2-, 3-, 4-, and 5-year nominal zero-coupon bonds, γ' is a vector of regression coefficients and f_t is a vector of the 1-, 2-, 3-, 4-, and 5-year annual forward rates and an intercept. We call this fitted value, $\hat{\gamma}' f_t$, the Cochrane-Piazzesi factor. They then run individual excess bond returns regressions of the form

$$rx_{t+1}^{(n)} = b_n (\hat{\gamma}' f_t) + \varepsilon_{t+1}^{(n)}, \quad (41)$$

where b_n is the maturity-specific regression coefficient. We replicate their methodology both in the historical sample and in the model simulated samples. The rightmost columns in Table 8 shows that the Cochrane-Piazzesi factor indeed is a strongly significant predictor of

future excess bond returns in the model as well as in the data. Consistent with the data, the R^2 's are decreasing in bond maturity, while the regression coefficient is increasing in bond maturity. While the evidence is strong in the model, it is stronger in the data. The R^2 's in the model are all around 4%, while in the data they are more than 30%! So, from these forecasting regressions the annual 5-year bond risk premium in the model has an implied standard deviation of 1.2%, while the corresponding number in the data is 3.4%. While the model does not replicate the full magnitude of the deviations from the expectations hypothesis as measured by these regressions, a standard deviation of 1.2% for a 5-year bond risk premium that is only 1.1% unconditionally, is highly economically significant (see, e.g., Kandel and Stambaugh, 1996).

6 Conclusion

This paper makes two main contributions. First, we present a two-good model where only one of the goods, the basic good, is exposed to an external unobservable habit. We show that the relative price between the two goods can be used to substitute out the unobservable habit, which leaves us with an, in principle, *observable* stochastic discount factor. This allows us to *estimate* the habit formation model and so we avoid the reverse engineering that usually is necessary when calibrating the dynamic behavior of the habit level (e.g., Campbell and Cochrane, 1999; Wachter, 2006). Using available disaggregate U.S. data on nondurables and services, we construct a basket of goods which reasonably can be labelled as luxury goods. The remaining goods in the nondurables and services group are labelled basic goods. We show that our definition of luxury good consumption is more volatile than and more correlated with equity returns than the standard nondurable and services aggregate consumption measure that is usually employed, consistent with the findings in Ait-Sahalia, Parker, and Yogo (2004). We show that of these reasons, the model can account for the unconditional equity premium puzzle, the risk-free rate puzzle, and the excess volatility puzzle. The model can also account for the unconditional level of bond risk premia observed in the data.

Second, we estimate the conditional volatility of the two risk factors, luxury good consumption and relative price growth, to operate at a business cycle and a lower-frequency, "generational" cycle, respectively. Thus, the price of risk and risk premium dynamics in the model are predicted to operate at both of these frequencies. The model predicts that the lower-frequency fluctuations in risk are quantitatively more important for risk premia than

the business cycle fluctuations. We indeed find evidence in the data for both a business cycle and a lower frequency component in the equity risk premium. Further, consistent with the model, the lower frequency component is quantitatively the more important one. The same dynamic behavior holds for bond risk premia in the model. Since the model predicts counter-cyclical bond risk premia, it also captures the failure of the expectations hypothesis. In particular, we show that the Cochrane and Piazzesi (2005) factor predicts future bond returns in the model at all maturities. The average simulated sample from the model does not, however, quantitatively generate as much predictability as that found in the data.

The estimated conditional volatility of the relative price of basic versus luxury goods predicts both excess bond and equity returns in the data, and it is significantly positively correlated with both the aggregate dividend yield and a measure of the aggregate consumption-wealth ratio (the *cay*-variable of Lettau and Ludvigson, 2001). The empirical relevance of this variable, which is estimated from macro-data only, is important evidence in favor of the two-good model advocated in this paper.

In future research, it will be useful to consider richer specifications for the exogenous state processes of the risk factors, aggregate dividends and inflation than those estimated in this model. This may improve the quantitative match of the conditional moments of the model with those estimated from the data.

References

- [1] Abel (1990), "Asset Prices under Habit Formation and Catching up with the Joneses", *American Economic Review* 80, 38 - 42
- [2] Ait-Sahalia, Parker and Yogo (2004), "Luxury Goods and the Equity Premium", *Journal of Finance* 59, 2959 - 3004.
- [3] Ang, A. and M. Piazzesi (2003), "A No-Arbitrage Vector Autoregression of Term Structure Dynamics with Macroeconomic and Latent Variables", *Journal of Monetary Economics*, 50, 4, 745-787.
- [4] Bansal and Yaron (2004), "Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles", *Journal of Finance* 54, 1481-1509.

- [5] Boldrin, Christiano and Fisher (2001), "Habit Persistence, Asset Returns, and the Business Cycle", *The American Economic Review*, Vol. 91, No. 1, pp. 149-166.
- [6] Bollerslev, T. "Modeling the Coherence in Short-Run Nominal Exchange Rates: A Multivariate Generalized ARCH Model." *Review of Economics and Statistics*, 1990, vol. 72, 498 - 505.
- [7] Bollerslev, T. and Wooldridge, J. M. "Quasi-Maximum Likelihood Estimation and Inference in Dynamic Models with Time-Varying Covariances." *Econometric Reviews*, 1992, vol. 11, 83 - 106.
- [8] Boudoukh, Jacob, Roni Michaely, and Matthew Richardson, and Michael Roberts, "On the Importance of Measuring Payout Yield: Implications for Empirical Asset Pricing", *Journal of Finance*, 2007, 62, 877 - 915.
- [9] Calvet and Fisher (2005), "Multifrequency News and Stock Returns", NBER Working Paper 11441
- [10] Campbell, J. (1999), "Asset Prices, Consumption, and the Business Cycle", Chapter 19 in John Taylor and Michael Woodford eds., *Handbook of Macroeconomics*, Amsterdam: North-Holland, 1999.
- [11] Campbell and Cochrane (1999), "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior", *Journal of Political Economy* 107, 205 - 251
- [12] Campbell and Shiller (1988), "The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors", *Review of Financial Studies* 58, 195 - 227
- [13] Campbell and Shiller (1991), "Yield spreads and interest rate movements: A bird's eye view", *Review of Economic Studies* 58, 495-514.
- [14] Chan and Kogan (2002), "Catching up with the Joneses: Heterogeneous Preferences and the Dynamics of Asset Prices", *Journal of Political Economy* 109
- [15] Chetty and Szeidl (2007), "Consumption Commitments and Risk Preferences", *Quarterly Journal of Economics* 122(2): 831-877.
- [16] Chetty and Szeidl (2005), "Consumption Commitments: Neoclassical Foundations for Habit Formation", Working paper, UC Berkeley.

- [17] Lawrence J. Christiano, Martin Eichenbaum and David Marshall (1991), "The Permanent Income Hypothesis Revisited ", *Econometrica*, Vol. 59, No. 2 (Mar., 1991), pp. 397-423.
- [18] Cochrane, J., and M. Piazzesi (2005), "Bond Risk Premia", *American Economic Review* Volume 95, Issue 1, March 2005, pp. 138-160.
- [19] Duffee, Greg (2006), "Term structure estimation without using latent factors," *Journal of Financial Economics* 79, 507-536.
- [20] Fama, E. and R. Bliss "The Information in Long-Maturity Forward Rates", *American Economic Review*, 1987, 77(4), pp. 680-92.
- [21] Fama and French (1989), "Business Conditions and Expected Returns on Stocks and Bonds", *Journal of Financial Economics* 25, 23 - 49
- [22] Fama (1991), "Efficient Capital Markets II", *Journal of Finance* 46, 1575 - 1617.
- [23] Gabaix, X. and D. Laibson (2002), "The 6D Bias and the Equity Premium Puzzle", with David Laibson, NBER Macroeconomics Annual, vol. 16, pp. 257-312.
- [24] Heaton (1995), "An Empirical Investigation of Asset Pricing with Temporally Dependent Preference Specifications," *Econometrica* 63, 681-717.
- [25] Jagannathan, Ravi, and Yong Wang. 2007. "Calendar Cycles, Infrequent Decisions and the Cross Section of Stock Returns," working paper, Northwestern University.
- [26] Jermann, U. (1998), "Asset Pricing in Production Economies", *Journal of Monetary Economics*, April, 257-275.
- [27] Kandel and Stambaugh, (1996), "On the Predictability of Stock Returns: An Asset-Allocation Perspective," *Journal of Finance*, 51, 385-424.
- [28] Lettau and Ludvigson (2001), "Consumption, Aggregate Wealth and Stock Returns", *Journal of Finance* 3, 815 - 849
- [29] Lettau and Ludvigson (2007), "Measuring and Modeling Variation in the Risk-Return Tradeoff", prepared for the *Handbook of Financial Econometrics*, editors Yacine Ait-Sahalia and Lars P. Hansen.

- [30] Lustig and Van Nieuwerburgh (2005), "Housing Collateral, Consumption Insurance and Risk Premia: an Empirical Perspective", *Journal of Finance* 60, pp. 1167-1219.
- [31] Lynch, A. (1996). "Decision frequency and synchronization across agents: Implications for aggregate consumption and equity returns". *Journal of Finance* 51(4):1479–1497.
- [32] Mankiw and Zeldes (1991), "The Consumption of Stockholders and Non-Stockholders," *Journal of Financial Economics* 29, 97-112.
- [33] Mehra and Prescott (1985), "The Equity Premium: A Puzzle", *Journal of Monetary Economics* 15, 145 - 161
- [34] Melino, Angelo and Alan X. Yang, "State-dependent preferences can explain the equity premium puzzle" *Review of Economic Dynamics* 6 (October 2003): 806-830.
- [35] Nelson, D. "Conditional Heteroskedasticity in Asset Returns: A New Approach." *Econometrica*, 1991, vol. 59, 347-370.
- [36] Newey, W., and D. McFadden. "Large Sample Estimation and Hypothesis Testing," in *Handbook of Econometrics* (Vol. 4), 1994, eds. R. Engle and D. McFadden, New York: Elsevier Science, pp. 2113-2245.
- [37] Parker, J. and C. Julliard (2005), "Consumption Risk and the Cross-Section of Expected Returns," *Journal of Political Economy*, 113(1), February 2005.
- [38] Piazzesi, Schneider and Tuzel (2006), "Housing, Consumption and Asset pricing", *Journal of Financial Economics* 83, pp. 531-569.
- [39] Shiller, R. (1981), "Do Stock Prices Move Too Much to be Justified by Subsequent Changes in Dividends?", *American Economic Review*, 71(3): 421–436.
- [40] Vissing-Jorgensen (2001), "Limited Stock Market Participation and the Elasticity of Intertemporal Substitution," *Journal of Political Economy*.
- [41] Wachter (2006), "A Consumption-Based Model of the Term Structure of Interest Rates", *Journal of Financial Economics* 79, 365-399.
- [42] Weil (1989), "The Equity Premium Puzzle and the Riskfree Rate Puzzle", *Journal of Monetary Economics* 24, 401–421

- [43] Wilcox (1992), "The Construction of U.S. Consumption Data: Some Facts and Their Implications for Empirical Work." *American Economic Review*, 82 (September), 922–941.
- [44] Working, H. (1960), "Note on the Correlation of First Differences of Averages in a Random Chain", *Econometrica*, Vol. 28, No. 4, pp. 916-918 .
- [45] Yogo (2006), "A Consumption-Based Explanation of Expected Stock Returns", *Journal of Finance* 61, 539-580

7 Appendix

7.1 Estimating the EGARCH Process

The system to be estimated is

$$\Delta l_t = a_l + b_l \sigma_{l,t}^2 + \sigma_{l,t} \varepsilon_{l,t} \quad (42)$$

$$\log \sigma_{l,t}^2 = \omega_l + \beta_{1,l} |\varepsilon_{l,t-1}| + \beta_{2,l} \varepsilon_{l,t-1} + \beta_{3,l} \log \sigma_{l,t-1}^2 \quad (43)$$

and

$$\Delta p_t = a_p + b_p \sigma_{p,t}^2 + \sigma_{p,t} \varepsilon_{p,t} \quad (44)$$

$$\log \sigma_{p,t}^2 = \omega_p + \beta_{1,p} |\varepsilon_{p,t-1}| + \beta_{2,p} \varepsilon_{p,t-1} + \beta_{3,p} \log \sigma_{p,t-1}^2 \quad (45)$$

where $E(\varepsilon_{l,t}, \varepsilon_{p,t}) = \rho$ and $\varepsilon_{l,t}, \varepsilon_{p,t} \sim N(0, 1)$.

The log-likelihood function conditional on the $t = 1$ variances is:

$$\mathcal{L}_T(\theta) = \sum_{t=1}^T l_t(\theta) = -\frac{1}{2} \sum_{t=1}^T \left(2 \ln 2\pi + \ln \sigma_{l,t}^2 + \ln \sigma_{p,t}^2 + \ln(1 - \rho^2) + \frac{\varepsilon_{l,t}^2 + \varepsilon_{p,t}^2 - 2\rho \varepsilon_{l,t} \varepsilon_{p,t}}{1 - \rho^2} \right), \quad (46)$$

where

$$\varepsilon_{l,t} = (\Delta l_t - a_l - b_l \sigma_{l,t}^2) / \sigma_{l,t}, \quad (47)$$

$$\varepsilon_{p,t} = (\Delta p_t - a_p - b_p \sigma_{p,t}^2) / \sigma_{p,t}, \quad (48)$$

and ρ is the correlation coefficient between $\varepsilon_{l,t}$ and $\varepsilon_{p,t}$. Note that a two-step estimation first over the parameters of each variable's univariate EGARCH process and then over the correlation parameter will lead to consistent estimates (see Newey and McFadden, 1994, for sufficient conditions for consistency and asymptotic normality of such estimates). Therefore, we first estimate the EGARCH model for each variable individually and subsequently compute the correlation between the residuals. We take the corresponding estimates as starting values for the full, joint likelihood estimation of equation (46).

7.1.1 The Score of the Constant Correlation EGARCH-in-mean Model

Let $y_t = [\Delta l_t \ \Delta p_t]'$ and θ the parameter vector. Define the conditional mean and variance of y_t as

$$E_t(y_t) \equiv \mu_t(\theta), \quad (49)$$

$$V_t(y_t) \equiv \Omega_t(\theta), \quad (50)$$

$$\varepsilon_t \equiv y_t - \mu_t(\theta), \quad t = 1, 2, \dots, T. \quad (51)$$

There are P parameters. The $1 \times P$ log-likelihood score function $s_t(\theta)$ is then given by:

$$\begin{aligned} s_t(\theta)' &\equiv \frac{\partial}{\partial \theta} l_t(\theta)' = \frac{\partial \mu_t(\theta)'}{\partial \theta} \Omega_t^{-1} \varepsilon_t(\theta) \\ &+ \frac{1}{2} \frac{\partial \Omega_t'}{\partial \theta} [\Omega_t^{-1}(\theta) \otimes \Omega_t^{-1}(\theta)] \text{vec} [\varepsilon_t(\theta) \varepsilon_t(\theta)' - \Omega_t(\theta)], \end{aligned} \quad (52)$$

where $\frac{\partial \mu_t(\theta)}{\partial \theta}$ is the $2 \times P$ derivative of $\mu_t(\theta)$ and $\frac{\partial \Omega_t}{\partial \theta}$ is the $2^2 \times P$ derivative of $\Omega_t(\theta)$.²⁰ The derivative of the a matrix A is defined here as the derivative of the vector $\text{vec}(A)$. In the case at hand, we have that $\theta \in \{a_l, b_l, a_p, b_p, \omega_l, \beta_{1l}, \beta_{2l}, \beta_{3l}, \omega_p, \beta_{1p}, \beta_{2p}, \beta_{3p}, \rho\}$. Further,

$$\mu_t(\theta) = \begin{bmatrix} a_l + b_l \sigma_{l,t}^2 \\ a_p + b_p \sigma_{p,t}^2 \end{bmatrix}, \quad (53)$$

and

$$\Omega_t(\theta) = \begin{bmatrix} \sigma_{l,t}^2 & \rho \sigma_{l,t} \sigma_{p,t} \\ \rho \sigma_{l,t} \sigma_{p,t} & \sigma_{p,t}^2 \end{bmatrix}, \quad (54)$$

²⁰The vec -operator denotes the vectorization of a matrix into a column vector by stacking the columns of the matrix on top of each other. The \otimes -operator denotes the Kronecker product.

where

$$\sigma_{l,t}^2 = \exp(\omega_l + \beta_{l,1} |\varepsilon_{l,t-1}| + \beta_{l,2} \varepsilon_{l,t-1} + \beta_{l,3} \ln \sigma_{l,t-1}^2), \quad (55)$$

$$\sigma_{p,t}^2 = \exp(\omega_p + \beta_{p,1} |\varepsilon_{p,t-1}| + \beta_{p,2} \varepsilon_{p,t-1} + \beta_{p,3} \ln \sigma_{p,t-1}^2). \quad (56)$$

7.2 Data Construction

This section explains how the luxury good consumption data and relative prices have been constructed. The National Income and Product Accounts (NIPA) tables are provided by the Bureau of Economic Analysis (BEA). Underlying the aggregate Personal Consumption Expenditures for durables, nondurables, and services are the "Details" tables, which we use to construct the basic good and luxury good consumption baskets, as well as their implicit price deflators. The model in this paper calls for nondurable and services consumption expenditures. Based on the three principles outlined in the main text, we rely on the subgroups Hotels and Motels, Airfare, Beauty Parlors and Health Clubs, and Private Flying, to construct a proxy for aggregate luxury good consumption. The remaining subgroups (for instance, items under food and clothing) are deemed as basic good consumption expenditures.²¹ The resulting proxies for basic good and luxury good consumption expenditures thus add up to the reported aggregate nondurable and services consumption expenditures. The data series goes back to 1959Q1 and the sample used in the paper is 1959Q1 – 2006Q4 (192 quarters).

We simply add the dollar expenditures on each of the four identified luxury groups to get dollar luxury good expenditures, $P_{L,t}^{\$} L_t$. In the model, basic goods are, without loss of generality, taken to be the numeraire:

$$C_t = P_t L_t + B_t, \quad (57)$$

where P_t is the price of the luxury good relative to the basic good, and C_t is total consumption

²¹There are some consumption subgroups that arguably could be classified as luxury goods that we have chosen to classify as basic goods. Casino gambling, for instance, is clearly discretionary. However, most of the income in this sector comes from the lower-end part of the market and also gambling addiction is something we do not want to capture. Brokerage expenditures is another subgroup of services consumption that could be considered luxury consumption. However, these expenses may be better thought of as investment. We therefore chose to not include this subgroup, although including this measure would have helped the model as brokerage expenses are highly correlated with stock market returns and very volatile.

expenditures in terms of the basic good (as opposed to \$'s). The model has a representative agent, so the consumption is per capita. As explained in the main text, we need measures of P_t and L_t to estimate the parameters of their assumed growth rate processes. We use the same methodology as NIPA in doing this (see Appendix 1 of their article "A Guide to the National Income and Product Accounts of the United States). That is,

$$\frac{P_{L,t}^{\$}}{P_{L,t-1}^{\$}} = \sqrt{\frac{\sum_i P_{i,t} Q_{i,t-1} \sum_i P_{i,t} Q_{i,t}}{\sum_i P_{i,t-1} Q_{i,t-1} \sum_i P_{i,t-1} Q_{i,t}}}, \quad (58)$$

and

$$\frac{L_t^*}{L_{t-1}^*} = \sqrt{\frac{\sum_i P_{i,t-1} Q_{i,t} \sum_i P_{i,t} Q_{i,t}}{\sum_i P_{i,t-1} Q_{i,t-1} \sum_i P_{i,t} Q_{i,t-1}}}, \quad (59)$$

where $P_{i,t}$ is the time t price given for each subgroup i and $Q_{i,t}$ is the time t quantity consumed of each subgroup i . The P 's and the Q 's are taken from the NIPA tables as discussed above. Note that L^* is the total consumption as reported and *not* per capita consumption, L_t . Reassuringly,

$$\frac{P_{L,t}^{\$} L_t^*}{P_{L,t-1}^{\$} L_{t-1}^*} = \sqrt{\frac{\sum_i P_{i,t} Q_{i,t-1} \sum_i P_{i,t} Q_{i,t}}{\sum_i P_{i,t-1} Q_{i,t-1} \sum_i P_{i,t-1} Q_{i,t}}} \sqrt{\frac{\sum_i P_{i,t-1} Q_{i,t} \sum_i P_{i,t} Q_{i,t}}{\sum_i P_{i,t-1} Q_{i,t-1} \sum_i P_{i,t} Q_{i,t-1}}} \quad (60)$$

$$= \frac{\sum_i P_{i,t} Q_{i,t}}{\sum_i P_{i,t-1} Q_{i,t-1}}. \quad (61)$$

Next, We need the price of basic goods in order to calculate the relative price that is consistent with the model. First, define basic good expenditures in the data implicitly using the relation

$$P_{L,t}^{\$} L_t^* + P_{B,t}^{\$} B_t = \text{consumption expenditures nondurables and services} = P_{C,t}^{\$} C_t, \quad (62)$$

where $P_{C,t}^{\$}$ and C_t are given in the NIPA consumption expenditure tables. Using this, the

price growth for basic goods is given by

$$\frac{P_{B,t}^{\$}}{P_{B,t-1}^{\$}} = \sqrt{\frac{(P_{C,t}^{\$}C_{t-1} - P_{L,t}^{\$}L_{t-1}^*) (P_{C,t}^{\$}C_t - P_{L,t}^{\$}L_t^*)}{(P_{C,t-1}^{\$}C_{t-1} - P_{L,t-1}^{\$}L_{t-1}^*) (P_{C,t-1}^{\$}C_t - P_{L,t-1}^{\$}L_t^*)}}, \quad (63)$$

$$= \sqrt{\frac{P_{C,t}^{\$}C_t - P_{L,t}^{\$}L_t^*}{P_{C,t-1}^{\$}C_{t-1} - P_{L,t-1}^{\$}L_{t-1}^*} \frac{P_{C,t}^{\$}}{P_{C,t-1}^{\$}} \frac{C_{t-1}}{C_t} \frac{1 - \frac{P_{C,t-1}^{\$}}{P_{C,t}^{\$}} \frac{P_{L,t}^{\$}}{P_{L,t-1}^{\$}} \frac{P_{L,t-1}^{\$}L_{t-1}^*}{P_{C,t-1}^{\$}C_{t-1}}}{1 - \frac{P_{C,t}^{\$}}{P_{C,t-1}^{\$}} \frac{P_{L,t-1}^{\$}}{P_{L,t}^{\$}} \frac{P_{L,t}^{\$}L_t^*}{P_{C,t}^{\$}C_t}}}. \quad (64)$$

Now, we can compute the relative price

$$\frac{P_t}{P_{t-1}} = \frac{P_{L,t}^{\$}}{P_{L,t-1}^{\$}} \frac{P_{B,t-1}^{\$}}{P_{B,t}^{\$}}, \quad (65)$$

and

$$\frac{L_t}{L_{t-1}} = \frac{L_t^*}{L_{t-1}^*} \frac{N_{t-1}}{N_t}, \quad (66)$$

where N_t is the population each quarter as given in the NIPA Personal Income account and L_t is the per capita luxury good consumption measure we use in the paper. Also, note that the inflation series used in the paper is $\frac{P_{B,t}^{\$}}{P_{B,t-1}^{\$}}$, consistent with the model. Since the measure of luxury good is only about 2% of total nondurable and services expenditures, this inflation measure is empirically very close to the usual deflator used for standard one-good models.