

Consumption Volatility Risk*

Oliver Boguth

Sauder School of Business
University of British Columbia

Lars-Alexander Kuehn

Tepper School of Business
Carnegie Mellon University

June 22, 2009

ABSTRACT

Empirically, the conditional volatility of aggregate consumption growth varies over time. While many papers test the consumption CAPM based on realized consumption growth, little is known about how the time-variation of consumption growth volatility affects asset prices. We show that in a model where (i) the agent has recursive preferences, (ii) the conditional first and second moments of consumption growth follow a Markov chain and (iii) the state of the economy is latent, the perception about condi-

I. Introduction

It is well known that the volatility of macroeconomic quantities, such as consumption and output, vary over time.¹ While many papers test the consumption CAPM based on realized consumption growth, e.g., Lettau and Ludvigson (2001b), Parker and Julliard (2005) and Yogo (2006), little is known about how the time-variation of consumption growth volatility affects both the cross-section and time-series of stock returns. The goal of this paper is to fill this void from a theoretical as well as empirical perspective. Specifically, we are interested in whether innovations to the conditional volatility of consumption growth are a priced risk factor.²

This research question poses several challenges: First, a natural candidate to model consumption volatility is the ARCH model proposed by Engle (1982) and its various generalizations. Asset pricing theory, however, states that only innovations are priced and in a GARCH model the volatility process has no separate innovations relative to the main process. In particular, Restoy and Weil (2004) show that a GARCH consumption model does not give rise to a volatility risk factor in an equilibrium model with Epstein and Zin (1989) utility. Second, while consumption growth rates are observable, the conditional volatility is latent and has to be estimated from the data. Lastly, aggregate consumption is measured with error thereby making statistical inference more difficult (Breedon, Gibbons, and Litzenberger (1989) and Wilcox (1992)).

In our model, which builds on Bansal and Yaron (2004), Lettau, Ludvigson, and Wachter (2008) and Kandel and Stambaugh (1991), the representative agent has recursive Epstein and Zin (1989) preferences and the conditional first and second moments of consumption growth follow independent two-state Markov chains. Recursive preferences imply that the agent cares not only about shocks to current consumption growth but also about changes to the distribution of future consumption growth, which in the model is driven by a persistent conditional mean and volatility. Since the state of the economy is unobservable, the agent uses Bayesian updating to form beliefs about the state, similar to David (1997) and Veronesi

¹For instance, refer to Cecchetti and Mark (1990), Kandel and Stambaugh (1990), Bonomo and Garcia (1994), Kim and Nelson (1999), or Whitelaw (2000).

²Other recent contributions testing the consumption CAPM include Campbell (1996), Ait-Sahalia, Parker, and Yogo (2004), Campbell and Vuolteenaho (2004), Bansal, Dittmar, and Lundblad (2005), Lustig and Nieuwerburgh (2005), and Jagannathan and Wang (2007).

(1999). Consequently, the agent's perception about the conditional first and second moment of consumption growth are priced. Given that both risk aversion and the elasticity of intertemporal substitution (EIS) are greater than one, changes about the perceived conditional mean carry a positive price of risk and changes of the perceived conditional consumption volatility a negative one.

The economic mechanism underlying our model is the following. When risk aversion and EIS are greater than unity, the intertemporal substitution effect dominates the wealth effect. As a result, the demand for the risky asset and thus the wealth-consumption ratio increases with expected consumption growth and decreases with consumption growth volatility. This effect also implies a positive price of risk for the first moment and negative one for the second moment of consumption growth. Intuitively, consider an asset that comoves negatively with future consumption growth. Its payoff is high (low) when investors learn that future consumption growth is low (high). Investors will demand a low return from this asset as it is a welcome insurance against future bad times. Similarly, consider an asset that comoves highly with future consumption volatility. This asset has high (low) payoffs when investors

have lower future returns. A volatility risk (VR) factor, which is the return of holding a long position in the value-weighted quintile of stocks with high volatility risk and a short position in low volatility risk, has an average return of -5% per year. Importantly, consumption volatility risk quintiles do not display variation in the average book-to-market characteristic. Nevertheless, the loadings of the 25 size and value portfolios of Fama and French (1992) on the VR factor monotonically decrease in the book-to-market ratio.

The main implication of the model is that consumption volatility is a priced risk factor. In order to test this hypothesis, we perform two stage regressions of excess returns on log consumption growth, changes in the perceived mean and volatility of consumption and the VR factor. Crucially, the coefficient on the innovation in the perceived consumption volatility and VR are negative implying that the representative agent has EIS greater than one. This finding contributes to a long standing debate in the literature on the magnitude of the EIS. Early evidence suggests that the EIS is smaller than one, e.g., Hall (1988) and Campbell and Mankiw (1989). More recently, Attanasio and Weber (1993), Vissing-Jorgensen (2002) and Vissing-Jorgensen and Attanasio (2003) find the opposite.

We also augment the CAPM and Fama-French 3-factor model with the VR factor. In particular, the VR factor shows up strongly and significant in addition to the market and the three Fama and French (1993) factors. Adding the VR factor to specifications that already contain HML, the overall fit of the regression as measured by the R^2 statistic improves only marginally. At the same time, replacing HML by VR has little effect on the goodness of the model. We thus conclude that HML and VR are substitutes in the cross-sectional pricing relation. But in contrast to HML, the volatility risk factor has a clear economic interpretation.

Another implication of our model is the predictability of the equity premium in the time-series. In states with low conditional mean or high conditional volatility of consumption growth, the model predicts a high equity premium when the representative agent has risk aversion and EIS greater than unity. We show in a predictive regression that innovations to consumption volatility are a significant and robust predictor of one-quarter ahead equity returns. The R^2 of this univariate predictive regression is almost 5% and a one standard deviation increase of the perceived consumption volatility results in a 1.8% rise of the quarterly equity premium. Both values are very close to the predictive power of the consumption-wealth ratio cay of Lettau and Ludvigson (2001a), the best known macroeconomic predictor of the

short horizon equity premium. In our model, changes in consumption volatility enter the pricing kernel only because they affect the consumption-wealth ratio. Thus, one might expect that measures of the consumption-wealth ratio such as `cay` already contain information about the volatility state. Empirically, this is not the case. Both variables are virtually uncorrelated and both remain strong and robust predictors in multivariate settings. The in-sample R^2 of regressions containing both changes in volatility and `cay` exceed 10%.

In the literature, it is common to measure consumption risk by using non-durable plus service consumption. This assumption is usually justified with a felicity function which is separable across goods. With Epstein-Zin utility, however, felicity can be separable across goods, but due to the time-nonseparability of the time-aggregator, other goods still matter for asset pricing because they enter the pricing kernel via the wealth-consumption ratio. The wealth-consumption ratio can be a function, for instance, of human capital (e.g. Jagannathan and Wang (1996), Lettau and Ludvigson (2001b), and Santos and Veronesi (2006)), durable goods (e.g. Yogo (2006)) or housing consumption (e.g. Piazzesi, Schneider, and Tuzel (2007)). If the wealth-consumption were observable, it would subsume all these variables. But the wealth-consumption ratio is unobservable.³ The contribution of this paper is to show that the conditional volatility of consumption growth is a significant determinant of the wealth-consumption ratio by documenting that it is priced in the cross-section and time-series after controlling for other factors.

Related Literature

Pindyck (1984) and Poterba and Summers (1986) are among the first to show that a decrease in prices is generally associated with an increase in future volatility, the so-called leverage or volatility feedback effect. Similarly, Campbell and Hentschel (1992) and Glosten, Jagannathan, and Runkle (1993) look at the relation between market returns and market volatility in the time-series. More recently, Ang, Hodrick, Xing, and Zhang (2006) use a nonparametric measure of market volatility, namely, the option implied volatility index (VIX), to show that innovations in aggregate market volatility carry a negative price of risk in the cross-section. Adrian and Rosenberg (2008) use a GARCH inspired model to decompose market volatility into short and long run components and show how each of the two components affects the

³One of the first papers which tries to estimate the wealth-consumption ratio is Lettau and Ludvigson (2001a). A more recent contribution is Lustig, Van Nieuwerburgh, and Verdelhan (2008).

cross-section of asset prices.

All of the papers mentioned above use some measure of stock market volatility, and can therefore be interpreted as extensions of a market based CAPM. We differ from the existing research in two important dimensions. First, we extract aggregate volatility from consumption data and not from financial data. While the quality of consumption data is considerably worse than the quality of financial data, we are able to robustly show that consumption volatility risk is priced. Second, we explicitly model the fact that conditional moments of consumption growth are unobservable and investors learn about them.

Considerable less research has been done on the pricing implications of volatility in a consumption-based model. Notable exceptions are Jacobs and Wang (2004) and Balduzzi and Yao (2007), who use survey data to estimate the variability of idiosyncratic consumption across households. They find that exposure to idiosyncratic consumption risk bears a negative risk premium for the 25 Fama-French portfolios. Parker and Julliard (2005) empirically measure a version of long-run risk as the covariance between one-period asset returns and long-horizon movements in the pricing kernel. Their ultimate consumption risk measure performs favorably in explaining the return differences of the 25 Fama-French portfolios. Similarly, Tedongap (2007) estimates conditional consumption volatility as a GARCH process and finds that value stocks covary more negatively with changes in consumption volatility over long horizons, thus requiring high average returns. In contrast to Tedongap (2007), we extract innovations to beliefs about consumption volatility, whereas a GARCH model does not allow that. Tedongap (2007) obtains significant results only at long horizons since GARCH models account for innovations to volatility only through realized data.

Motivated by the long-run risk model of Bansal and Yaron (2004), there exist important papers which study the relation between aggregate volatility and prices. Notably, Bansal, Khatchatrian, and Yaron (2005) find that the conditional consumption volatility predicts valuation ratios. Drechsler and Yaron (2008) extend the long-run risk model for jumps in consumption volatility. Their model generates a variance premium and return predictability which is consistent with the data. Bansal and Shaliastovich (2008) find evidence that measures of investors uncertainty about their estimate of future growth contain information about large moves in returns at frequencies of about 18 months. They explain this regularity with a recursive-utility based model in which investors learn about latent expected consumption

growth from signals with time-varying precision. Bollerslev, Tauchen, and Zhou (2008) study the asset pricing implication when the variance of stochastic volatility is stochastic.⁴

More closely related are Calvet and Fisher (2007) and Lettau, Ludvigson, and Wachter (2008). Calvet and Fisher (2007) study the asset pricing implications of multi-fractal Markov switching in a recursive preference model at the aggregate level. Similarly, Lettau, Ludvigson, and Wachter (2008) estimate a Markov model with learning to show that the decline in consumption volatility—also referred to as the “Great Moderation”—can explain the high observed stock market returns in the 1990s and the following decline in equity risk premium. We extend their work by studying the cross-section and time-series of returns.

The remainder of the paper is organized as follows: In Section 2, we derive the asset pricing implication of a recursive preference model where the agent does not observe the state of the economy. This section motivates our empirical analysis of Sections 3–5. In Section 3, we test whether consumption growth and its conditional moments forecast returns in the cross-section. We run Fama-MacBeth regressions and form portfolios based on consumption volatility loadings. In Section 4, we test whether consumption growth and its conditional moments as well as the VR factor are priced risk factors. Section 5 contains time-series predictability tests and Section 6 concludes.

II. Model

In this section, we derive the asset pricing implications of a model where the representative agent has recursive preferences and the state of the economy is unobservable. A crucial implication of recursive preferences is that the agent cares not only about shocks to current consumption growth but also about news regarding the distribution of future consumption growth. In our model, future consumption growth is influenced by time-variation of the conditional mean and volatility of consumption growth which is unobservable to the agent. This latent nature implies the agent forms beliefs about the conditional first and second moments of consumption growth and, most importantly, changes in the perceived first and second moments of consumption growth are priced.

⁴Other papers building on the long-run risk framework of Bansal and Yaron (2004) include Bhamra, Kuehn, and Strebulaev (2007), Hansen, Heaton, and Li (2008) and Bansal, Dittmar, and Kiku (2009)

A. Consumption

We assume that the conditional first and second moments of consumption growth follow a Markov chain. Specifically, let Δc_{t+1} denote log consumption growth, μ_t its conditional expectation and σ_t its conditional volatility, then log consumption growth follows

$$\Delta c_{t+1} = \mu_t + \sigma_t \epsilon_{t+1} \quad \epsilon_t \sim \mathcal{N}(0; 1) \quad (1)$$

with iid innovations ϵ_t . For tractability in the empirical estimation, we assume two states for the mean and two for the volatility which are denoted by $\mu_t \in \{l; h\}$ and $\sigma_t \in \{l; h\}$. The conditional first and second moments of consumption growth follow Markov chains with transition matrix P^μ and P^σ , respectively, given by

$$P^\mu = \begin{bmatrix} p_{ll}^\mu & 1 - p_{ll}^\mu \\ 1 - p_{hh}^\mu & p_{hh}^\mu \end{bmatrix} \quad P^\sigma = \begin{bmatrix} p_{ll}^\sigma & 1 - p_{ll}^\sigma \\ 1 - p_{hh}^\sigma & p_{hh}^\sigma \end{bmatrix} \quad (2)$$

To keep the number of parameters to be estimated manageable, we impose that mean and volatility states switch independently. Thus, the joint transition matrix is the product of the marginal probabilities for mean and volatility and the 16-element matrix can be fully characterized by 4 parameters. Importantly, the assumption of independent switching probabilities does not imply that mean and volatility or the beliefs thereof are independent. Since we assume two drift and two volatility states, there are four states in total, $\{(\mu_t, \sigma_t) \in \{l; h\} \times \{l; h\}\}$, denoted by $S_t = 1, \dots, 4$. Our specification follows Kandel and Stambaugh (1990), Kim and Nelson (1999), and Lettau, Ludvigson, and Wachter (2008).

In contrast to Bansal and Yaron (2004) and Kandel and Stambaugh (1991), we assume that the representative agent does not observe the state of the economy. Instead, she must infer it from observable consumption data. This assumption ensures that the empirical exercise is in line with the model. The inferences at date t about the underlying state is captured by the posterior probability of being in each state based on the available data Y_t . Let $\pi_{t+1|t}$ denote the posterior belief vector of size 4×1 :

$$\pi_{t+1|t} = P' \frac{P' \pi_{t|t-1} \odot \epsilon_t}{1'(P' \pi_{t|t-1} \odot \epsilon_t)} \quad (3)$$

where

$$\mathbf{f}_t = \begin{bmatrix} f(\Delta c_t | \mathbf{y}_{t-1} = \mathbf{l}; \mathbf{y}_{t-1} = \mathbf{l}; Y_{t-1}) \\ f(\Delta c_t | \mathbf{y}_{t-1} = \mathbf{l}; \mathbf{y}_{t-1} = \mathbf{h}; Y_{t-1}) \\ f(\Delta c_t | \mathbf{y}_{t-1} = \mathbf{h}; \mathbf{y}_{t-1} = \mathbf{l}; Y_{t-1}) \\ f(\Delta c_t | \mathbf{y}_{t-1} = \mathbf{h}; \mathbf{y}_{t-1} = \mathbf{h}; Y_{t-1}) \end{bmatrix}$$

is a vector of Gaussian likelihood functions and $\mathbf{P} = \mathbf{P}^\mu \otimes \mathbf{P}^\sigma$ is the transition matrix.

B. Recursive Utility

The representative household maximizes recursive utility over consumption following Kreps and Porteus (1978), Epstein and Zin (1989), and Weil (1989):

$$U_t = \left\{ (1 - \beta) C_t^\rho + \left(\mathbb{E}_t[U_{t+1}^{1-\gamma}] \right)^{\rho/(1-\gamma)} \right\}^{1/\rho} \quad (4)$$

where C_t denotes consumption, $\beta \in (0, 1)$ the rate of time preference, $\rho = 1 - 1/\psi$ and γ the elasticity of intertemporal substitution (EIS), and γ relative risk aversion. Implicit in the utility function (4) is a constant elasticity of substitution (CES) time aggregator and CES power utility certainty equivalent.

Epstein-Zin preferences provide a separation of the elasticity of intertemporal substitution and relative risk aversion. These two concepts are inversely related when the agent has power utility. Intuitively, the EIS measures the agents willingness to postpone consumption over time, a notion well-defined even under certainty. Relative risk aversion measures the agents aversion to atemporal risk across states.⁵

We know from Epstein and Zin (1989) that the Euler equation for any return $R_{i,t+1}$ can be stated as

$$\mathbb{E}_t \left[\theta \left(\frac{P C_{t+1} + 1}{P C_t} \right)^{-(1-\theta)} \left(\frac{C_{t+1}}{C_t} \right)^{-\gamma} R_{i,t+1} \right] = 1 \quad (5)$$

where $\theta = \frac{1-\gamma}{1-1/\psi}$ and $P C_t = P_t C_t$ denotes the wealth-consumption ratio. For the empirical exercise, it is useful to study the log-linearized pricing kernel. A log-linear approximation of the pricing kernel implicit in (5) is

$$m_{t+1} \approx (\ln \theta - (1 - \theta)k_0) - \Delta c_{t+1} - (1 - \theta)(p c_{t+1} - k_1 p c_t) \quad (6)$$

⁵Recursive preferences also imply preference for early or late resolution of uncertainty. This feature, however, is not important for this paper since the agent cannot choose between consumption lotteries which differ in the timing of the resolution of uncertainty.

where $\text{pc}_t = \ln(P_t/C_t)$ denotes the log wealth-consumption ratio and k_0, k_1 are constants. The value of k_1 is given by $k_1 = \overline{PC} / (\overline{PC} - 1) > 1$ and \overline{PC} is the mean wealth-consumption ratio. The Epstein-Zin pricing kernel is thus a function of consumption growth and the level of the log wealth-consumption ratio. Alternatively, the pricing kernel can approximately be stated in terms of changes of the log wealth-consumption ratio if k_1 is close to one. Lustig, Van Nieuwerburgh, and Verdelhan (2008) estimate the unconditional quarterly wealth-consumption ratio to be almost 351 implying that $k_1 = 1.003$. Consequently, the log pricing kernel can be closely approximated by:

$$m_{t+1} \approx (\ln - (1 - \gamma)k_0) - \Delta c_{t+1} - (1 - \gamma)\Delta \text{pc}_{t+1} \quad (7)$$

The log pricing kernel (7) implies that excess returns are determined as covariance between returns and log consumption growth as well as the covariance between returns and changes of the log wealth-consumption ratio:

$$\mathbb{E}_t[\mathbf{R}_{i,t+1}^e] \approx -\text{Cov}_t(\mathbf{R}_{i,t+1}, m_{t+1}) = -\text{Cov}_t(\mathbf{R}_{i,t+1}, \Delta c_{t+1}) + (1 - \gamma)\text{Cov}_t(\mathbf{R}_{i,t+1}, \Delta \text{pc}_{t+1}) \quad (8)$$

In an endowment model which is solely driven by iid shocks, the wealth-consumption ratio is constant. In our model, however, the first and second moments of consumption growth follow a Markov chain. The unobservability of the state implies that the agent's beliefs characterize the state of the economy. Consequently, the wealth-consumption ratio is a function of the agent's beliefs about the state of the economy, i.e., $PC_t = PC(\cdot |_{t+1|t})$.

For the pricing of the return on the consumption claim, Euler equation (5) simplifies to

$$PC_t^\theta = \mathbb{E}_t \left[\theta (PC_{t+1} + 1)^\theta \left(\frac{C_{t+1}}{C_t} \right)^{1-\gamma} \right] \quad (9)$$

Based on the law of iterated expectations, equation (9) implies that

$$PC_t^\theta = \sum_{i=1}^4 \pi_{t+1|t}(i) PC_{t,i}^\theta \quad (10)$$

where $\pi_{t+1|t}(i)$ is i -the element of $\pi_{t+1|t}$ and

$$PC_{t,i}^\theta = \mathbb{E} \left[\theta (PC_{t+1} + 1)^\theta \left(\frac{C_{t+1}}{C_t} \right)^{1-\gamma} \middle| s_{t+1} = i; \pi_{t+1|t} \right] \quad (11)$$

Equation (10) says that the agent forms a belief-weighted average of the state- and belief-conditioned wealth-consumption ratios (11).

In order to study how the wealth-consumption ratio changes with beliefs about the state, we further define the posterior belief that the mean state is high tomorrow by

$$b_{\mu,t+1|t} = P(\mu_{t+1} = \mu_h | Y_t) = b_{t+1|t}(3) + b_{t+1|t}(4) \quad (12)$$

and the posterior belief that the volatility state is high tomorrow by

$$b_{\sigma,t+1|t} = P(\sigma_{t+1} = \sigma_h | Y_t) = b_{t+1|t}(2) + b_{t+1|t}(4) \quad (13)$$

The univariate effects of changing beliefs about the volatility (mean) while holding the mean (volatility) constant can locally be approximated. Given the volatility, changes of the log wealth-consumption ratio are

$$\Delta pc_{t+1} \approx (b_{\mu,t+2|t+1} - b_{\mu,t+1|t}) \frac{1}{b_{\mu,t+1|t} PC_{\mu=\mu_h,\sigma}^\theta + (1 - b_{\mu,t+1|t}) PC_{\mu=\mu_l,\sigma}^\theta} PC_{\mu=\mu_h,\sigma}^\theta - PC_{\mu=\mu_l,\sigma}^\theta \quad (14)$$

where $PC_{\mu=\mu_h,\sigma}$ denotes the wealth-consumption ratio when expected consumption growth is high and the consumption volatility is constant, a similar definition applies for the other wealth-consumption ratios. Analogous, given the mean, changes of the log wealth-consumption ratio are

$$\Delta pc_{t+1} \approx (b_{\sigma,t+2|t+1} - b_{\sigma,t+1|t}) \frac{1}{b_{\sigma,t+1|t} PC_{\mu,\sigma=\sigma_h}^\theta + (1 - b_{\sigma,t+1|t}) PC_{\mu,\sigma=\sigma_l}^\theta} PC_{\mu,\sigma=\sigma_h}^\theta - PC_{\mu,\sigma=\sigma_l}^\theta \quad (15)$$

Changes in the log wealth-consumption ratio are thus locally proportional to changes in beliefs. From an empirical asset pricing perspective, this finding implies that changes in beliefs are priced in the time-series and cross-section since they affect the wealth-consumption ratio, according to Equation (8).

This implication does not necessarily follow from an equilibrium model where the conditional consumption volatility follows a GARCH process. In a GARCH model, the conditional volatility is a function of lagged volatility and lagged squared residuals of the consumption process. Thus, a GARCH process is not driven by separate innovations relative to the consumption process. Consequently, Restoy (1991) and Restoy and Weil (2004) have shown that

to its covariance with consumption growth. Volatility, which affects the wealth-consumption ratio, therefore can have pricing implications as it determines the loading on the consumption growth factor, but it does not give rise to a second priced risk factor. Restoy and Weil continue to say: "This result embodies the fundamental insight that, for a GARCH(1,1) process, returns are only able to predict future conditional means of consumption growth but carry no information about the future conditional variances".

C. Estimation

To estimate the model, we obtain data on quarterly per capita consumption from the NIPA tables as the sum of nondurables and services. In accordance with the observation that the consumption behavior in the United States in the years following World War II is systematically different from later years, we restrict our time-series from the first quarter of 1955 until the fourth quarter of 2007. The choice of 1955 allows sufficient consumption observations to ensure that the impact of prior beliefs on posteriors has vanished by the time we start the portfolio analysis in 1964.

The resulting parameter estimates of the Markov chain are reported in Table I, Panels A and B. Expected consumption growth is always positive and about twice as high in the high state relative to the low state ($\mu_l = 0.375\%$, $\mu_h = 0.748\%$ quarterly). State-conditioned consumption volatilities are $\sigma_l = 0.211\%$ and $\sigma_h = 0.463\%$. The probability of remaining in a given regime for the mean is 0.935 in the low state and 0.912 in the high state. The volatility regimes are somewhat more persistent, with probabilities of 0.977 and 0.978, respectively. Our estimates differ from the ones presented by Lettau, Ludvigson, and Wachter (2008), who estimate volatility in both states to be more persistent (0.991 and 0.994). In addition to small differences in the time-series, their consumption measure is total consumption, while we use the common definition of consumption as nondurables plus services.

We assume independent switching in mean and volatility states. This assumption greatly reduces the number of parameters to be estimated and thus improves estimation precision. Yet this assumption does not appear to be a significant restriction of consumption data. If the true (unobservable) correlation between regime switches were very different from zero, we would expect to see a large correlation in our estimated posterior beliefs. Panel C of Table I shows that this correlation is less than 15% which is small enough to suggest that the

assumption of independence is not a contradiction of the data.

Figure 1 shows the filtered beliefs for the regimes. Panel A depicts the belief dynamics for mean consumption growth $\mathbf{b}_{\mu,t+1|t}$ and Panel B for the standard deviation $\mathbf{b}_{\sigma,t+1|t}$. These graphs visually confirm that the mean regimes are less persistent than the standard deviation states. In particular, the parameter estimates for the Markov chain imply that mean states last for 3.4 years whereas volatility states last for 11.2 years on average. Further, a decline in consumption volatility from the 1990s onwards, as pointed out by Kim and Nelson (1999), is easily observable. Importantly, there is significant variation in beliefs about the volatility state in addition to decline in the early 1990s.

D. Implications

Based on the parameters estimated from consumption data, we solve the model numerically to study its properties. In the following, we are interested in how the perception about the first and second moments of consumption growth affect the wealth-consumption ratio. To this end, we define expected consumption growth and expected consumption volatility as a belief weighted average:

$$\hat{\mu}_{t+1|t} = \mathbf{b}_{\mu,t+1|t} \mathbf{h} + (1 - \mathbf{b}_{\mu,t+1|t}) \mathbf{l} \quad \hat{\sigma}_{t+1|t} = \mathbf{b}_{\sigma,t+1|t} \mathbf{h} + (1 - \mathbf{b}_{\sigma,t+1|t}) \mathbf{l} \quad (16)$$

In Figures 2 and 3, we plot the wealth-consumption ratio as a function of the perceived conditional first $\hat{\mu}_{t+1|t}$ (left graph) and second $\hat{\sigma}_{t+1|t}$ (right graph) moments of consumption growth for the benchmark calibration of Table I when the agent has a high EIS of 1.5 (Figure 2) and low EIS of 0.5 (Figure 3). We further calibrate the model to a quarterly rate of time preference of $\beta = 0.995$ and risk aversion of $\gamma = 30$. Risk aversion of 30 seems unrealistically high. This section, however, is meant to yield qualitative guidance and not quantitative results.

To gain a better understanding of the economics, it is convenient to recall the Gordon growth model. Under the assumption that discount and growth rates are constant, the Gordon growth model states that the wealth-consumption ratio is negatively related to the risk-free rate r_f and risk premium r_E and positively to the growth rate g , i.e., $PC = 1/(r_f + r_E - g)$.

In Figure 2, the wealth-consumption ratio is increasing in the perceived mean and decreasing in the perceived volatility of consumption growth when the EIS equals 1.5. In Figure 3, the opposite is true when the EIS equals 0.5. Interestingly, both, Figures 2 and 3 suggest

that the log wealth-consumption ratio is approximately affine in the perceived first and second moment of consumption growth for our parameterization, confirming the validity of Equations (14) and (15).

The sign change in the slope of the wealth-consumption ratio with respect to expected consumption growth is driven by two opposing effects. On the one hand, a higher perception about the growth rate increases the wealth-consumption ratio as in the Gordon growth model. On the other hand, in equilibrium, an increase in expected consumption growth also raises the risk-free rate since the riskless asset becomes less attractive relative to the risky asset. This second effect lowers the wealth-consumption ratio. When the EIS is greater than unity, the first effect (intertemporal substitution effect) dominates the second effect (wealth effect). As a result, the demand for the risky asset and thus the wealth-consumption ratio rises with the perceived expected growth rate of consumption.

Similarly, the sign change in the slope of the wealth-consumption ratio with respect to expected consumption growth volatility (Figure 2 versus 3) is also driven by two opposing effects. On the one hand, a higher perceived conditional consumption volatility increases the risk premium which lowers the wealth-consumption ratio as in the Gordon growth model. On the other hand, in equilibrium, an increase in expected consumption growth volatility also reduces the risk-free rate since the riskless asset becomes more attractive relative to the risky asset. This effect increases the wealth-consumption ratio. If $\gamma > 1$, the first effect dominates the second when the EIS is greater than one.

In order to test this intuition in the cross-section of returns, it is convenient to restate the fundamental asset pricing equation (8) in terms of betas:⁷

$$\mathbb{E}_t[\mathbf{R}_{i,t+1}^e] \approx \beta_{c,t}^i \mathbf{c}_{c,t} + \beta_{pc,t}^i \mathbf{p}_{pc,t} \approx \beta_{c,t}^i \mathbf{c}_{c,t} + \beta_{\mu,t}^i \boldsymbol{\mu}_{\mu,t} + \beta_{\sigma,t}^i \boldsymbol{\sigma}_{\sigma,t} \quad (17)$$

where $\beta_{c,t}^i, \beta_{pc,t}^i, \beta_{\mu,t}^i, \beta_{\sigma,t}^i$ denote risk loadings of asset i at date t with respect to consumption growth, the wealth-consumption ratio and the conditional first and second moment of consumption growth and $\mathbf{c}_{c,t}, \mathbf{p}_{pc,t}, \boldsymbol{\mu}_{\mu,t}, \boldsymbol{\sigma}_{\sigma,t}$ are the respective market prices of risk. The main cross-sectional implications of the model are the following. Assuming that both risk aversion and EIS are greater than one, the agent requires lower expected excess returns for stock

⁷Equation (10) states that variations in the wealth-consumption ratio depend on the beliefs about four states, three of which are linearly independent. In an exact implementation of the model, the wealth-consumption ratio is thus a nonlinear function of three variables. We reduce it to two in order to achieve a more meaningful economic interpretation for mean and volatility states. Note, however, that the empirical results are not affected by this assumption.

which load less (low betas) on expected consumption growth and more (high betas) on future consumption growth volatility.

The economic intuition is straightforward. Equations (8) and (17) imply that the market prices of the conditional first and second moments of consumption growth are their respective conditional variances, each multiplied by a constant.⁸ The sign of both constants depends on $(1 - \theta)$, coming from Equation (8), and on the slope of the wealth-consumption ratio with respect to the conditional mean and volatility of consumption growth. Assuming both risk aversion and EIS greater than one, $(1 - \theta)$ is positive and the wealth-consumption ratio is upward-sloping in the conditional mean and downward-sloping in the conditional volatility of consumption growth (Figure 2). Consequently, the market price of expected consumption growth is positive, i.e. $\mu_{\mu,t} > 0$, and the market price of conditional consumption volatility is negative, i.e. $\sigma_{\sigma,t} < 0$.

Moreover, the positive sign of $\mu_{\mu,t}$ and negative one of $\sigma_{\sigma,t}$ also implies that investors require higher compensation for holding stocks which load strongly (high beta) on expected consumption growth and less compensation for stocks which load strongly (high beta) on consumption growth volatility. Intuitively, assets, which comove negatively with future consumption growth, have high payoffs when investors learn that future consumption growth is low. These assets thus provide insurance against future bad times. Similarly, assets, which comove highly with future consumption volatility, have high payoffs when investors learn that future consumption is very volatile. These assets serve as insurance against uncertain times and thus have lower required returns.

III. Cross-Sectional Return Predictability

The goal of this section is to demonstrate that loadings on the perceived (filtered) conditional consumption volatility forecast returns. To this end, we first run quarterly time-series regressions to obtain loadings on risk factors. Next, we test using both Fama-MacBeth regressions and portfolio sorts whether these risk loadings forecast returns. Our main finding is that innovations in consumption volatility risk is a strong and robust predictor of future returns,

⁸Intuitively, the log wealth-consumption is approximately a line in the perceived first and second moments of consumption growth which expressed in changes can be stated as $\log pc_t \approx A \mu_t + B \sigma_t$. Substituting this expression into (8), one obtains (17) and $\lambda_{\mu,t} = A(1 - \theta)\text{Var}_t(\mu_t)/\mathbb{E}_t[M_{t+1}]$ and $\lambda_{\sigma,t} = B(1 - \theta)\text{Var}_t(\sigma_t)/\mathbb{E}_t[M_{t+1}]$.

while exposure to consumption growth and changes in expected consumption growth risk do not help to explain the cross-section of asset prices.

A. Data

Our sample consists of all common stocks (`shrcd` = 10 or 11) on CRSP that are traded on the NYSE or AMEX (`exchcd` = 1 or 2). While the results are generally robust to the inclusion of NASDAQ stocks, this restriction mitigates concerns that only a small fraction of total market capitalization has a large impact on the portfolio analysis. To obtain valid risk measurements for a given quarter, the asset is required to have at least 60 months of prior data and at least 16 out of 20 valid quarterly returns. Since we use size and book-to-market ratio as characteristics, we require market capitalization to be available in December that occurs 7 to 18 months prior to the test month as well as book value of equity from COMPUSTAT in the corresponding year. The choice of the long delay is motivated by the portfolio formation strategies in Fama and French (1992), who want to ensure that the variables are publicly available when they are used in the study. Due to limited availability of book values in earlier years, we begin the empirical exercise in January 1964. The first time-series regression to estimate risk loadings thus covers the time span from 1959 to 1963. We end our analysis in December 2007.

B. Consumption Volatility Risk and Stock Returns

Our first set of empirical results is based on time-series regressions of individual securities onto log consumption growth, the perceived conditional mean and volatility of consumption growth as well as the excess return on the market. This specification is a generalization of the pricing restriction (8), where changes in the wealth-consumption ratio can be linearly approximated as shown in Equations (14) and (15). To be conservative, we also include the market return in the regression to control for time effects so that consumption-based betas purely load on cross-sectional differences in returns.

In particular, for each security, we estimate factor loadings in each quarter t^* using the previous 20 quarterly observations from

$$R_t^i - R_t^f = \beta_{t^*}^i + \beta_{M,t^*}^i (R_{M,t} - R_t^f) + \beta_{c,t^*}^i \Delta c_t + \beta_{\mu,t^*}^i \Delta \mu_t + \beta_{\sigma,t^*}^i \Delta \sigma_t + \beta_t^i \quad (18)$$

where R_t^f denotes the risk-free rate and $R_{M,t}$ the market return for $t \in \{t^* - 19; t^*\}$. Further,

Δ_{t-1} and Δ_t are innovations in the perceived conditional moments of consumption growth. These are defined as the differences in believed moments before and after consumption is realized:

$$\Delta_{t-1} = \hat{c}_{t+1|t-1} - \hat{c}_{t|t-1} \quad \Delta_t = \hat{c}_{t+1|t} - \hat{c}_{t|t-1} \quad (19)$$

where \hat{c}_{t-1} and \hat{c}_t are the perceived conditional first and second moments of consumption growth, respectively, defined in Equation (16).

The estimated parameters from Equation (18) allow us to evaluate the cross-sectional predictive power of these loadings in two different ways. First, we use cross-sectional regressions as in Fama and MacBeth (1973) to investigate if the factor loadings help to predict cross-sectional variation in returns. Second, we form portfolios based on the estimated risk exposures and analyze their properties in the time-series.

C. Fama-MacBeth Regressions

We now investigate the predictive power of the estimated loadings from model (18) by cross-sectionally regressing the returns in month $s+1$ of each asset onto its latest available risk loadings as well as size and value characteristics:

$$\begin{aligned} R_{s+1}^i = & \alpha_{0,s+1} + \alpha_{1,s+1} \hat{M}_{t^*}^i + \alpha_{2,s+1} \hat{c}_{t^*}^i + \alpha_{3,s+1} \hat{\mu}_{t^*}^i + \alpha_{4,s+1} \hat{\sigma}_{t^*}^i \\ & + \alpha_{5,s+1} ME_{t^*}^i + \alpha_{6,s+1} BM_{t^*}^i + \alpha_{7,s+1}^i \end{aligned} \quad (20)$$

The explanatory variables are normalized each quarter so they are centered around zero with unit variance. Equation (20) states that each set of three monthly regressions in one quarter will share the same predictor variables. For example, the returns in each of the months April, May, and June are regressed onto the risk loadings estimated from the window ending in the first quarter of the same year. We are interested whether the factor loadings have any predictive power for the cross-sectional variation of returns or, equivalently, whether $\alpha_{k,s}; k = 1, \dots, 6$ are on average different from zero.

The results of the Fama-MacBeth regressions are presented in Table II. Model specifications I-IV present univariate effects of each risk loading. Consistent with prior research (Mankiw and Shapiro (1986), Lettau and Ludvigson (2001b)), an asset's contemporaneous short horizon loadings on the market and consumption growth do not help to predict cross-sectional differences in returns. The average coefficient $\bar{\alpha}_1$ and $\bar{\alpha}_2$ are small and insignificant.

The same holds for innovations in the belief about the expected growth rate, $\bar{\gamma}_3$. Exposure to consumption volatility risk, however, as measured by $\bar{\gamma}_4$, shows up strongly negative and significant. Stocks that comove highly with changes in consumption volatility underperform their peers in the future. Specification V is the full model. Now, both the loading on consumption growth and consumption volatility risk are significant. In regression VI, we add two characteristics known to predict stock returns, namely, the market capitalization ($\bar{\gamma}_5$) and the ratio of book value of equity to market value ($\bar{\gamma}_6$), to confirm that the predictive power of consumption volatility is not already captured by these predictors. The absolute value of the point estimate $\bar{\gamma}_4$ is slightly reduced by the addition of the two characteristics, but it remains significant.

What do these findings mean? The novel implications of our model are that beliefs about mean and volatility states of consumption growth are priced sources of risk. As a result, exposure to these sources should be associated with a spread in future returns. The sign of the risk premium associated with each of these two factors depends on preference parameters. In the case where both risk aversion and EIS are greater than unity, the model predicts that returns are positively related to $\mu_{t,t}$ and negatively to $\sigma_{t,t}$. An EIS lower than one yields the opposite predictions. We do not find convincing evidence that exposure to fluctuations in expected consumption growth predicts returns. Yet the coefficient on consumption volatility is strongly negative. This finding is consistent with the model only if risk aversion and EIS of the representative agent is greater than one.

There are two potential problems in our analysis. First, when comparing the influence of the predictor variables, it is important to note the possibility of an error-in-variable problem. The first four variables are estimates from first stage regressions and are thus noisy. Moreover, while researchers often treat Δc_t as observable, the consumption time-series actually is measured with significant noise (Breedon, Gibbons, and Litzenberger (1989) and Wilcox (1992)). Both $\Delta \mu_t$ and $\Delta \sigma_t$ are estimates and themselves depend on the imposed model for consumption growth dynamics. The t-statistics reported are obtained from standard OLS theory and are likely based on standard errors that are too small. However, the necessary adjustment would require to keep track of standard errors across all stages of estimation and is thus not feasible in this case. In the next section, we use a different methodology to confirm that our results are robust and not due to misrepresented standard errors.

Second, imposing a Markov model on consumption growth leads to a collinearity problem in the first stage regression (18). Rational Bayesian updating leads investors to lower their belief about the expected consumption growth rate whenever the realized growth rate is lower than their prior belief. Since the distance between the two mean states is small relative to the total variation of consumption growth, this leads to a high correlation between ΔC_t and $\Delta \pi_t$. We estimate it to be 0.60 over the entire sample, and it ranges between 0.39 and 0.93 in 5-year subperiods. To avoid this collinearity, we omit either variable in the first stage regression (18). In untabulated results we confirm that the cross-sectional impact of volatility beliefs remains robust.

To reduce noise in measured consumption, we also estimate a very basic consumption tracking portfolio similar to Vassalou (2003) which also attenuates collinearity. To this end, we regress consumption growth on the market return, HML and SMB. The fitted values estimated from this regression can be interpreted as returns on a financial portfolio. This additional step does not only reduce the collinearity identified above, but also mitigates the concerns about measurement error in consumption data. A projection of consumption growth onto the space of financial payoffs eliminates possible noise in consumption data which is orthogonal to financial markets. At the same time, valuable information that is contained in the consumption series but not in financial data might be lost. In unreported results, we confirm previous findings.

D. Portfolio Sorts

An alternative approach to utilize the estimated risk loadings from regression (18) is to group the estimates cross-sectionally and form portfolios. This approach has several important advantages relative to Fama-McBeth regressions. First, the error-in-variable problem that led to underestimated standard errors in the regression approach now leads to conservatism in the statistical inferences. When variables are measured with noise, the portfolio sorts will be less accurate as some stocks will be assigned to the wrong portfolio. Under the assumption of cross-sectional predictive power, this leads to smaller return differences across portfolios. Since the statistical inference is based solely on portfolio returns, the measurement error ultimately leads to a decrease in statistical significance.

Second, forming portfolios allows to identify a non-linear relation between risk and ex-

pected excess returns. While the pricing kernel in equation (7) indicates that expected excess returns are approximately log-linear in the wealth-consumption ratio, the wealth-consumption ratio itself is not linear in beliefs about the mean or volatility states. Lastly, the portfolio approach results in a time-series of returns, which allows a further analysis of the relation between this strategy and other known risk factors.

At the end of each quarter, we sort all stocks in our sample into portfolios based on their estimated risk loadings from the time-series regression (18). Table III reports the average returns of equally-weighted (EW) and value-weighted (VW) quintiles as well as a long-short strategy that each month invests \$1 into quintile 5 (high risk) and sells \$1 of quintile 1 (low risk).

In Panel A, portfolios are formed based on loadings with respect to the market. Neither weighting scheme results in a measurable return dispersion across portfolios. A similar result obtains by forming portfolio based on exposure to consumptions growth $\hat{c}_{c,t}^i$ (Panel B) and changes in beliefs about expected consumption growth $\hat{\mu}_{\mu,t}^i$ (Panel C). In contrast, consumption volatility risk $\hat{\sigma}_{\sigma,t}^i$ shows up strongly negatively (Panel D). An equally-weighted strategy results in a return of the long-short portfolio of -0.32% monthly. The value weighted return is even larger (in absolute value) with -0.43% per month or in excess of -5% annually. In both cases, average returns across quintiles are monotonically decreasing. The large difference in returns is thus not driven by extreme observations in quintiles 1 and 5. Overall, the results in this table confirm the findings from the cross-sectional regressions in Table II.

Cross-sectional differences in returns might not be surprising if consumption volatility betas covary with other variables known to predict returns. Crucially, Table IV shows that this is not the case for the firm characteristics size and book-to-market. In Panel A, we again report average returns for each consumption volatility exposure quintile and its average beta. Panel B reports firm characteristics for each portfolio. Since market capitalization is non-stationary, and the value characteristic varies dramatically over time, we compute size- and value deciles for each stock at each month and take the average over these deciles within each portfolio. The table reports time-series means of portfolio characteristics. For market equity, we observe that the two extreme quintiles are composed of somewhat smaller than average stocks. This effect often shows up when ranking stocks by a covariance measure. Returns of small stocks are on average more volatile and risk estimates are therefore more likely to be

very large or very small. However, there is no difference in size rank between quintile 1 and 5. Most importantly, there is no variation in the book-to-market ratio across portfolios. Thus, consumption risk portfolios do not load on firm characteristics which are known to predict future returns.

A number of so-called anomalies are confined to a small subsets of stocks, often just to small companies or illiquid stocks (e.g. Fama and French (2008), Avramov, Chordia, Jostova, and Philipov (2007)). In Tables V and VI, stocks are independently sorted into three portfolios based on $\hat{\sigma}_{\sigma,t}^i$, and into two portfolios based on market capitalization (Table V) or book-to-market ratio (Table VI). The number of portfolios for each variable follows Fama and French (1993) and trades off the desire to obtain sufficient dispersion along each dimension while keeping the number of stocks in each portfolio large enough to minimize idiosyncratic risk. The bivariate sort in Table V shows that consumption volatility risk is consistently present and strong for both equal and value-weighted strategies with return differences ranging from -0.20% to -0.35% monthly. Interestingly, the effect is stronger for big than for small companies since returns of smaller stocks have a larger idiosyncratic component, and thus the risk estimates from the first stage regression are less precise. With these findings, there is no reason to believe that the predictive power of consumption volatility risk is associated with possible mispricing or slow information diffusion in small stocks. Similarly, Table VI confirms that consumption volatility risk is also present within book-to-market groups.

IV. Consumption Volatility Risk Pricing

Building on the findings of the previous section, we now investigate the pricing implications of beliefs about consumption moments cross-sectionally using the 25 Fama and French (1992) portfolios. These portfolios have been shown to challenge the single factor CAPM. Fama and French (1993) propose two additional factors that help to explain the return differences: size (SMB) and value (HML). While a convincing, unified economic interpretation for these factors is still outstanding, from an econometric view, the three factor model has been proven very successful.

We show that changes in beliefs about consumption volatility carry a negative price of risk, while changes in beliefs about the mean state do not contribute to explaining the cross-section of returns. To circumvent possible econometric problems related to multicollinearity in the

independent variables, we also form a long-short portfolio based on consumption volatility risk (VR) and demonstrate that it shows up strongly and significantly as a priced factor in cross-sectional regressions. While the VR portfolio only modestly correlates with HML, both factors are substitutes in the pricing relation. This evidence provides an economic interpretation for the risk associated with the HML factor.

A. Factor Pricing with Consumption Data

Equation (17) states that, in a log-linear approximation, expected excess returns depend on consumption growth and the perceived first and second moments of consumption growth. We evaluate the performance of our model in two stages. First, for each 25 Fama-French portfolio, we obtain risk loadings from the time-series regression

$$R_t^i - R_t^f = \alpha_i + \beta_c^i \Delta c_t + \beta_\mu^i \Delta \mu_t + \beta_\sigma^i \Delta \sigma_t + \epsilon_{it} \quad i = 1, \dots, 25 \quad (21)$$

In the second stage, we estimate the prices of risk by a cross-sectional regression of returns onto the loadings from the first stage.

Results from the second stage regression are summarized in Table VII. For each factor, the table reports point estimates for the prices of risk and associated t-statistics, which are adjusted for estimation error in the first stage as proposed by Shanken (1992) and are robust to heteroscedasticity and autocorrelation as in Newey and West (1987) with 4 quarterly lags. In addition, the following regression statistics are shown: The second stage R^2 and adjusted R^2 as well as the model J-test (χ^2 statistic) with its associated p-value (in percent). Return observations are at a quarterly frequency and the factors used are log consumption growth (Δc_t), changes in beliefs about the conditional mean of consumption growth ($\Delta \mu_t$), as well as changes in beliefs about consumption growth volatility ($\Delta \sigma_t$). A fourth factor, which is the return of a long-short portfolio that buys assets with high consumption volatility risk and sells assets with low consumption volatility risk, is also considered and denoted by VR.

Regression I shows results for the standard consumption CAPM. Confirming prior research, the market price of consumption risk is insignificant and an R^2 of about 8% indicates that the C-CAPM performs very poorly in pricing the chosen set of test assets. Regressions II - IV show the incremental effects of adding beliefs about conditional moments of consumption growth. Even though results in the previous section have shown that exposure to $\Delta \mu_t$ does not provide any return predictability, it is possible that $\Delta \mu_t$ carries a contemporaneous risk

premium. This is ruled out by regression II. The estimate for the price of risk on Δ_t is zero. In contrast, regression III shows that consumption volatility is a significant factor in the cross-section and, as the previous section suggests, the estimate for the price of volatility risk is negative.

Regression IV reports the full three factor model (21). Similar to the previous findings, both Δc_t and Δ_t are insignificant. The price of volatility risk, Δ_t , is negative, but lacks convincing statistical significance. This specification, however, suffers from a multicollinearity problem. While all explanatory variables are statistically insignificant, the model cannot be rejected, and the associated p-value of 96 percent is overwhelmingly large. An F-test of the joint significance confirms this intuition. The hypothesis that first stage risk loadings for Δc_t , Δ_t , and Δ_t are jointly equal to zero in the second stage is rejected at any significance level.⁹

While this evidence presents strong indication of multicollinearity, econometric theory is mostly silent about how to deal with it. We mitigate this concern by forming a consumption volatility risk (VR) portfolio as a proxy for Δ_t . The VR factor is a zero investment strategy that is long in the value-weighted quintile with the highest exposure and short in the value-weighted quintile with the lowest exposure to innovations in beliefs about consumption volatility as measured by $\hat{\sigma}_{\sigma,t}^i$ in Table III, Panel D. The estimated risk loadings are obtained from 20-quarter rolling regressions that end before the portfolio formation and are thus, conditional on consumption dynamics, publicly available information. We choose the univariate $\hat{\sigma}_{\sigma,t}^i$ dispersion as basis for the factor since the consumption volatility exposure is unrelated to size and value characteristics, as documented in Table IV.¹⁰

Regressions V and VI show a significantly negative price of risk for the VR factor, while beliefs about the mean consumption growth continue to be insignificant. The price for a unit of VR risk is estimated to be about 6% per quarter, or 25% annually. The specifications with VR perform well. The consumption CAPM augmented with the VR factor (Model VI) achieves second stage R^2 in excess of 70 percent, compared to 8 percent for the consumption

⁹Other indicators of multicollinearity are the Variance Inflation Factor (VIF) or the condition number of the matrix to be inverted ($X'X$). As a rule of thumb, multicollinearity is often considered problematic if the VIF is greater than 5 or if the condition number is greater than 20. In the given problem, both number are close to the corresponding values. The VIF is 4.25 and the condition number of the matrix ($X'X$) in the second stage is 16.

¹⁰We do not form a factor based on loadings on expected consumption growth (μ_t) since the spread between the high and low quintile is on average close to zero. Consequently, theory predicts that its market price of risk should be zero too.

CAPM.

The predictions of our theory in Section II depend on the preference parameters of the representative agent. Under the common assumption of risk aversion greater than unity, an EIS greater than one yields a negative price of risk for consumption volatility. While prior research often finds a negative price of risk for market volatility (Ang, Hodrick, Xing, and Zhang (2006), Adrian and Rosenberg (2008)), only a general equilibrium consumption-based model allows us to draw conclusions about preference parameters. The estimated prices of risk for both Δ_t and its mimicking VR portfolio are significantly negative, thus, suggesting an EIS greater than unity for the representative agent.

The second stage pricing errors are depicted in Figure 4. Each panel plots average quarterly excess returns against the model predicted excess returns for a given set of explanatory variables. If the model correctly prices assets and there are no errors induced from estimation or small sample size, all asset returns should line up exactly on the diagonal line. The top left graph depicts the consumption CAPM (regression I in Table VII). Visually, this graph confirms that the consumption CAPM does not work well in pricing the 25 Fama-French portfolios. While the portfolios vary drastically in their average returns, the model predicted returns are all very close together, resulting in a narrow cloud. In the top right graph, we augment the consumption CAPM with Δ_t (regression III in Table VII). Adding Δ_t as pricing factor helps to achieve more variation in predicted returns, but the pricing errors remain fairly large. In the two bottom graphs, we depict the full model with Δ_t (bottom left, regression IV), and the VR factor (bottom right, regression V). Both graphs confirm that in the full model, either with Δ_t or its mimicking portfolio, pricing errors are small and loadings on risk factors successfully explain average excess returns.

B. Factor Pricing with Portfolio Returns

To relate the pricing implications of consumption volatility risk to the existing literature, we now study market based rather than consumption based models. In particular, we are interested whether our VR factor is a substitute for Fama-French’s HML factor, thereby providing a macroeconomic interpretation for HML. Even though VR is independent of book-to-market characteristics and comoves only modestly with HML, we find that substituting HML by VR in the Fama-French three factor model results in similar pricing and leaves

pricing errors unaffected.

Summary statistics for the VR portfolio are given in Table VIII. The VR portfolio has a mean return of -0.43% and a standard deviation of 3.25% per month. Its standard deviation is lower than the market volatility, but comparable to the ones of the Fama-French factors. The monthly Sharpe ratio (in absolute value) of 0.13 is larger in magnitude than the Sharpe ratio of SMB (0.08) and close to the Sharpe ratio of HML (0.15). The correlation matrix of the pricing factors (Panel B) shows that the VR portfolio returns are uncorrelated with the market and SMB. Importantly, the correlation with the HML factor is moderate at -0.23 even though the VR portfolio is neutral with respect to the book-to-market characteristic (see Panel B of Table IV). Overall, the correlations of VR with all other factors are smaller than the pairwise correlations between the Fama French factors. Parameter estimates from a time-series regression of the VR factor onto the other factors are reported in Panel C. The CAPM (regression II) does not explain the returns of the VR portfolio. In regression III, the Fama French factors attenuate the estimated intercept $\hat{\alpha}$ slightly towards zero, but it remains large and significant. This attenuation is solely driven by HML and both the market and SMB have insignificant coefficients.

Using the same econometric methods as in the previous section, we now obtain risk loadings for each of the 25 Fama-French portfolios from time-series regressions. In the second stage, we estimate the prices of risk by a cross-sectional regression of returns onto the loadings from the first stage. Table IX reports factor loadings from first stage regressions of excess returns on the market excess return ($R_{M,t}^E$) and the VR factor

$$R_t^i - R_t^f = \alpha_i + \beta_M^i R_{M,t}^E + \beta_{VR}^i R_{VR,t} + \epsilon_t^i \quad i = 1; \dots; 25 \quad (22)$$

The risk estimates for both factors vary considerably across portfolios. Panel A confirms previous findings about market factor loadings. Along the size dimension, market betas decrease in size. While this is generally consistent with a risk based explanation of the size effect, the dispersion in betas is not sufficient to explain the large dispersion in returns. Along the value sorted portfolios, risk estimates actually decrease in the book-to-market ratio, while returns increase. This well known finding challenges a risk based explanation and contradicts the CAPM.

Risk exposures to the volatility risk factor are shown in Panel B. There is little variation in risk loadings β_{VR}^i across size portfolios. Small stocks have an average loading of -0.02 ,

and the loading monotonically decreases to -0.09 for large stocks. The dispersion in the risk loadings along the value dimension is much larger and decreases monotonically from 0.08 for growth stocks to -0.14 for value stocks. A low risk exposure is consistent with high expected returns for value stocks since the price of VR risk is negative. Loading on the VR factor therefore suggest a risk based explanation of the value anomaly.

Results from second stage regressions are reported in Table X. For each factor, the table presents point estimates for prices of risk. The associated t-statistics are based on standard errors that are Shanken (1992) and Newey and West (1987) adjusted. In addition, the following regression statistics are shown: The second stage R^2 and adjusted R^2 as well as the model J-test (χ^2 statistic) with its associated p-value (in percent). Regressions I and IV show the results for the benchmark models, the market CAPM (I) and the Fama-French model (IV). The CAPM does a very poor job in explaining the cross-section of returns. The point estimate for the market risk premium is negative and the regression R^2 is only 14%. The three factor model reduces the pricing errors significantly and yields an R^2 of 77%. The estimated market risk premium remains negative and the model is still rejected as indicated by the high χ^2 statistic.

The remaining regressions show various combinations of the Fama-French factors with VR. In all specifications, the estimates for the price of a unit VR risk are significant and negative, ranging from -1.21% to -2.84% monthly. To gauge the economic impact of VR on the cross-section of returns, we multiply average VR loadings of the value quintiles with the market price of VR risk. Panel B of Table IX shows that the average risk loadings monotonically decrease from 0.08 to -0.14 along the book-to-market characteristic, yielding a differential of -0.22 . Assuming a conservative price of VR risk (regression VI), this translates into an annual expected return differences between value and growth stocks of 3.8% .

In regression III, the factors are the market portfolio and VR as in Equation (22). This specification yields considerable improvements over the one factor market model and results in a similar fit relative to a model that contains the market and HML (regression VII). Interestingly, although VR is based on consumption data, a three factor model based on the market, SMB and VR (regression V) produces an R^2 of 80% which is slightly higher than the R^2 of the Fama-French model. Augmenting the Fama-French three-factor model with VR as a fourth factor (regression VI) does not lead to an improvement in the model's ability to

price the cross-section. In summary, replacing HML by VR does not deteriorate the model’s performance, while including both VR and HML as factors does not seem to improve the model fit. Hence, HML and VR are substitutes in cross-sectional pricing for the 25 Fama French portfolios. In contrast to HML, however, the consumption volatility risk portfolio has a clear economic interpretation.

Adrian and Rosenberg (2008) perform a similar analysis. They decompose stock market volatility into two components, which differ in persistence, and estimate them with a GARCH inspired model. In contrast, our VR portfolio is based on a Markov model for low-frequency consumption data. Interestingly, their short-run volatility component has similar pricing implications as VR, whereas their long-run component performs worse than VR. However, the persistence of their short-run volatility component is 0.327 for daily data while our consumption volatility regimes last on average for several years. The VR factor thus has a much different and macroeconomically more meaningful interpretation.

Figure 5 replicates Figure 4 for market based pricing models. The top left graph depicts the CAPM (regression I in Table X). The remaining graphs show the CAPM augmented with the volatility risk factor (top right graph, regression III), the Fama-French three factor model (bottom left graph, regression IV), and a three factor model that uses VR instead of HML (bottom right graph, regression V). Visually, these graphs confirm that simply adding VR to the market factor improves the model fit dramatically. Both the Fama-French model and the three factor VR model, though statistically still rejected, seem to price the 25 portfolios equally well.

V. Time Series Predictability

In the previous sections, we showed that loadings on consumption growth volatility predict returns cross-sectionally and consumption growth volatility is a priced risk factor. The model also predicts that the first and second moments of consumption growth forecast aggregate returns in the time-series. As explained in Section II, the model implies a negative relation between expected excess returns and expected consumption growth and a positive relation between expected excess returns and consumption growth volatility when both risk aversion and EIS are greater than unity. Noting that the wealth-consumption ratio is inversely related to expected excess returns, this effect can be seen in Figure 2. The opposite holds when risk

aversion and EIS are smaller than unity (see Figure 3).

We find that changes in consumption growth volatility are a strong and robust predictor of short horizon market returns whereas expected consumption growth is not a significant predictor variable. Specifically, higher consumption growth volatility is associated with higher future excess returns. This positive relation is consistent with previous results where loadings on consumption growth volatility are negatively related to returns. In a univariate regression of future returns on the consumption growth volatility, the regression R^2 is of similar magnitude as the one obtained in a univariate regression using the consumption-wealth ratio variable *cay* of Lettau and Ludvigson (2001a).

Table XI reports time-series regressions of quarterly excess market returns onto lagged predictor variables. Regressions I-III represent standard benchmark models. As predictor variables, we use the Treasury bill rate relative to its recent average as proposed by Hodrick (1992), the term spread, the default spread, and the dividend yield.¹¹ All of those have been shown to predict stock returns at various horizons.¹² Lettau and Ludvigson (2001a) use the household budget constraint to motivate the variable *cay* and show that it works exceptionally well at short horizon forecasts.¹³

In regressions IV and V, we study the predictive power of our two consumption state variables. Similar to the cross-sectional results in the previous sections, we find that beliefs about expected consumption growth do not predict stock returns, while changes in the volatility beliefs show up economically and statistically significant and yield an regression R^2 of almost 5% in a univariate regression. The R^2 of *cay* in the univariate regression II is slightly higher at about 6%. The economic impact of consumption volatility risk is large. A one standard deviation increase in Δ_t results in an increase in the expected risk premium of 1.8% quarterly.¹⁴ Regressions VI to VIII demonstrate that the marginal impact of Δ_t remains strong and significant even after controlling for all other predictors including *cay*. This is surprising since in our model, changes in consumption volatility enter the pricing kernel only

¹¹Data are from FRED. The relative Treasury bill rate is the yield on a 90 day T-bill less its past 12 months moving average. The term spread is the difference between yields of long (10 year) and short (1 year) government bonds, and the default spread is the yield difference between Baa rated and Aaa rated corporate bonds. The dividend yield is computed from CRSP as the ratio of gross cum-dividend index returns to gross ex-dividend returns.

¹²See, for example Roze (1984), Campbell and Shiller (1988), Fama and Schwert (1977), Keim and Stambaugh (1986), Campbell (1987), and Campbell and Thompson (2008).

¹³The variable *cay* was downloaded from Martin Lettau's homepage on Oct. 18, 2008.

¹⁴Note that σ_t has a standard deviation of around 0.0003.

because they affect the consumption-wealth ratio. Thus, one might expect that measures of the consumption-wealth ratio such as `cay` already contain information about the volatility state. In the data, however, this is not the case suggesting that `cay` is an imperfect measure of the wealth-consumption ratio.

It is well known that parameter estimates and t-statistics are potentially biased in predictive regressions. Persistence in stock returns, caused e.g. by using overlapping observations, leads to biased standard errors, which is examined in Hodrick (1992). In our setup, we do not use overlapping observations and the autocorrelation in market returns is very small. Another bias arises when the predictor variable is persistent and its innovations are correlated with future returns, as discussed in Stambaugh (1999), Lewellen (2004), Boudoukh, Richardson, and Whitelaw (2006) and Ang and Bekaert (2007). Especially when price ratios are used as predictors, this bias shows up strongly and conventional tests will reject the null hypothesis too frequently. To gain a better understanding, consider the following setup

$$\begin{aligned} r_t &= \alpha + \beta x_{t-1} + \varepsilon_t^r \\ x_t &= \gamma + \delta x_{t-1} + \varepsilon_t^x \end{aligned}$$

where r_t denotes returns and x_t a predictor variable. Lewellen (2004) shows that estimates are biased by $(\hat{\beta} - \beta)$ where $\hat{\beta} = \text{Cov}(\varepsilon_t^r; \varepsilon_t^x) / \text{Var}(\varepsilon_t^x)$ when ε_t^r is correlated with x_t . When the dividend yield is used as predictor, for instance, an increase in price leads to a positive realized return as well as a decrease in the dividend yield. Consequently, ε_t^r is correlated with x_t . Lewellen (2004) reports an auto-correlation of 0.997 and $\text{Corr}(\varepsilon_t^r; \varepsilon_t^x) = -0.96$ for the dividend yield as predictor, invalidating standard estimates and tests. For the variable Δ_t , this bias is less of a concern since it is not a price scaled variable. For our one period forecasts, we estimate $\text{Corr}(\Delta_t; \Delta_{t-1}) = 0.006$ and $\text{Corr}(\Delta_t; \varepsilon_t^x) = 0.032$, which is too small to bias statistical inference.

We acknowledge that the predictive results presented have limitations. First, they are in sample results. We leave evaluating the out of sample power for future work. Second, there is a look-ahead bias in Δ_t . In estimating the Markov chain for consumption growth, beliefs are updated according to Bayes' rule and therefore are not forward looking. The parameter estimates, however, are obtained by maximum likelihood, employing the full sample. In particular, investors in the early sample period know of the possibility that at one point in the future consumption volatility might switch to a state much lower than what had been

experienced in the past. This is similar to the critique by Brennan and Xia (2005), who point out that estimating cay over the entire sample induces a look ahead bias and a simple linear time trend would work as well as cay . Their criticism does not apply to our results since we use changes in beliefs as predictor which do not have a trend. Third, aggregate consumption data is not publicly available at the end of a quarter. Instead, initial estimates are published within the following month and they are subject to revisions for up to three years. Hence, we cannot conclude that it is possible to implement our predictability results in practice. Yet we succeed in identifying a new source of aggregate risk.

VI. Conclusion

When consumption growth is not i.i.d. over time and the representative household has recursive preferences, the wealth-consumption ratio is time-varying and enters the pricing kernel as a second factor (Epstein and Zin (1989), Weil (1989)). We generalize Bansal and Yaron (2004) to account for the latent nature of the conditional first and second moments of consumption growth. In the model, we identify innovations in beliefs about the conditional mean and volatility of consumption growth as two state variables that affect the wealth-consumption ratio and thus asset prices.

To test these predictions, we estimate a Markov model with two states for the conditional mean and two states for the conditional volatility of consumption growth, as in Kandel and Stambaugh (1990) and Lettau, Ludvigson, and Wachter (2008). Using the estimated beliefs from the Markov model, we empirically test the pricing implications for the cross-section and time-series of stock returns. We find that an asset's exposure to changes in beliefs about consumption volatility significantly forecast returns, while exposure to changes in beliefs about expected consumption growth do not.

In cross-sectional pricing tests using the 25 Fama-French size and value portfolios, both, changes in volatility and a portfolio, which is long assets with high volatility exposure and short assets with low volatility exposure, show up as priced factors. The effects are robust to a variety of pricing models, including augmented versions of the consumption CAPM, the market CAPM, and the Fama-French three factor model. The comparison of pricing errors across different model specifications indicates that the volatility risk factor has similar pricing implications to the value factor HML, even though, the volatility risk factor is neutral with

respect to the value characteristic.

In time-series tests, we find that innovations to beliefs about the volatility state forecast the equity premium. A one standard deviation increase in perceived volatility is followed by an increase of the equity returns of 1.8% quarterly. In a univariate regression, changes about perceived consumption volatility achieve an R^2 of about 5% which is comparable to predictive power of Lettau and Ludvigson (2001a) *cay* variable. The signs of the coefficients on consumption volatility in time-series and cross-sectional regressions indicate that the representative agent has risk aversion and elasticity of intertemporal substitution greater than unity.

Appendix

A. Numerical Solution

Since the price-consumption ratio is only a function of beliefs, i.e., $\text{PC}_t = \text{PC}(\mathbf{s}_{t+1|t})$, the price-consumption ratio solves:

$$\text{PC}(\mathbf{s}_{t+1|t}) = \left(\mathbb{E}_t \left[\theta (\text{PC}(\mathbf{s}_{t+2|t+1}) + 1)^\theta \left(\frac{\mathbf{C}_{t+1}}{\mathbf{C}_t} \right)^{1-\gamma} \right] \right)^{1/\theta}$$

The right-hand side can be simplified to

$$\text{PC}(\mathbf{s}_{t+1|t}) = \left(\sum_{i=1}^4 \mathbf{s}_{t+1|t}(\mathbf{i}) \mathbb{E} \left[\theta (\text{PC}(\mathbf{s}_{t+2|t+1}) + 1)^\theta (e^{\mu_i + \sigma_i \epsilon_{t+1}})^{1-\gamma} \middle| \mathbf{s}_{t+1} = \mathbf{i} \right] \right)^{1/\theta}$$

where $\mathbf{s}_{t+1|t}(\mathbf{i})$ is \mathbf{i} -the element of $\mathbf{s}_{t+1|t}$. We solve this equation as a fixed-point in the price-consumption. The grid for the belief state-vector has increments of size 0.025 and the expectation is approximated using Gauss-Hermite quadrature with 10 nodes. Linear interpolation is used for off-grid beliefs.

References

- Adrian, Tobias, and Joshua Rosenberg, 2008, Stock returns and volatility: Pricing the short-run and long-run components of market risk, *Journal of Finance* 63.
- Aït-Sahalia, Yacine, Jonathan A. Parker, and Motohiro Yogo, 2004, Luxury goods and the equity premium, *Journal of Finance* 59, 2959{3004.
- Ang, Andrew, and Geert Bekaert, 2007, Stock return predictability: Is it there?, *Review of Financial Studies* 20, 651{707.
- Ang, Andrew, Robert J. Hodrick, Yuhang H. Xing, and Xiaoyan Y. Zhang, 2006, The cross-section of volatility and expected returns, *Journal of Finance* 61, 259{299.
- Attanasio, O. P., and G. Weber, 1993, Consumption growth, the interest-rate and aggregation, *Review of Economic Studies* 60, 631{649.
- Avramov, Doron, Tarun Chordia, Gergana Jostova, and Alexander Philipov, 2007, Momentum and credit rating, *Journal of Finance* 62, 2503{2520.
- Balduzzi, Pierluigi, and Tong Yao, 2007, Testing heterogeneous-agent models: an alternative aggregation approach, *Journal of Monetary Economics* 54, 369{412.
- Bansal, Ravi, Robert F. Dittmar, and Dana Kiku, 2009, Cointegration and consumption risks in asset returns, Forthcoming, *Review of Financial Studies*.
- Bansal, Ravi, Robert F. Dittmar, and Christian T. Lundblad, 2005, Consumption, dividends, and the cross section of equity returns, *Journal of Finance* 60, 1639{1672.
- Bansal, Ravi, Varoujan Khatchatrian, and Amir Yaron, 2005, Interpretable asset markets?, *European Economic Review* 49, 531{560.
- Bansal, Ravi, and Ivan Shaliastovich, 2008, Confidence risk and asset prices, Working paper.
- Bansal, Ravi, and Amir Yaron, 2004, Risks for the long run: A potential resolution of asset pricing puzzles, *Journal of Finance* 59, 1481{1509.
- Bhamra, Harjoat S., Lars-Alexander Kuehn, and Ilya A. Strebulaev, 2007, The levered equity risk premium and credit spreads: A unified framework, University of British Columbia Working Paper.
- Bollerslev, Tim, Georg Tauchen, and Hao Zhou, 2008, Expected stock returns and variance risk premia, Working Paper.
- Bonomo, Marco, and Rene Garcia, 1994, Can a well- tted equilibrium asset-pricing model produce mean reversion, *Journal of Applied Econometrics* 9, 19{29.
- Boudoukh, Jacob, Matthew Richardson, and Robert F. Whitelaw, 2006, The myth of long-horizon predictability, Forthcoming, *Review of Financial Studies*.
- Breeden, Douglas T., Michael R. Gibbons, and R. H. Litzenberger, 1989, Empirical tests of the consumption-oriented capm, *Journal of Finance* 44, 231{262.
- Brennan, Michael J., and Yihong Xia, 2005, tay's as good as cay, *Finance Research Letters* 2, 1{14.
- Calvet, Laurent, and Adlai Fisher, 2007, Multifrequency news and stock returns, *Journal of Financial Economics* 86, 178{212.

- Campbell, John Y., 1987, Stock returns and the term structure, *Journal of Financial Economics* 18, 373{399.
- , 1996, Understanding risk and return, *Journal of Political Economy* 104, 298{345.
- , and Ludger Hentschel, 1992, No news is good-news - an asymmetric model of changing volatility in stock returns, *Journal of Financial Economics* 31, 281{318.
- Campbell, John Y., and Gregory Mankiw, 1989, Consumption, income, and interest rates: Reinterpreting the time series evidence, in O. Blanchard Fisher, and S., ed.: *National Bureau Economic Research Macroeconomics Annual: 1989*. pp. 185{216 (MIT Press: Cambridge, MA).
- Campbell, John Y., and Robert Y. Shiller, 1988, The dividend-price ratio and expectations of future dividends and discount rates, *Review of Financial Studies* 1, 195{228.
- Campbell, John Y., and Samuel B. Thompson, 2008, Predicting excess stock returns out of sample: Can anything beat the historical average?, *Review of Financial Studies* 21, 1509{1531.
- Campbell, John Y., and Tuomo Vuolteenaho, 2004, Bad beta, good beta, *American Economic Review* 94, 1249{1275.
- Cecchetti, Stephen G., and Nelson C. Mark, 1990, Evaluating empirical tests of asset pricing models: Alternative interpretations, *American Economic Review* 80, 48{51.
- David, Alexander, 1997, Fluctuating confidence in stock markets: Implications for returns and volatility, *Journal of Financial and Quantitative Analysis* 32, 427{462.
- Drechsler, Itamar, and Amir Yaron, 2008, What's volatility got to do with it, Working Paper.
- Engle, Robert F., 1982, Autoregressive conditional heteroscedasticity with estimates of the variance of united kingdom inflation, *Econometrica* 50, 987{1007.
- Epstein, Larry, and Stanley Zin, 1989, Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework, *Econometrica* 57, 937{969.
- Fama, Eugene F., and Kenneth R. French, 1992, The cross-section of expected stock returns, *Journal of Finance* 47, 427{465.
- , 1993, Common risk-factors in the returns on stocks and bonds, *Journal of Financial Economics* 33, 3{56.
- , 2008, Dissecting anomalies, *Journal of Finance* 63, 1653{1678.
- Fama, Eugene F., and James D. MacBeth, 1973, Risk, return, and equilibrium - empirical tests, *Journal of Political Economy* 81, 607{636.
- Fama, Eugene F., and G. William Schwert, 1977, Asset returns and inflation, *Journal of Financial Economics* 5, 115{146.
- Glosten, Lawrence R., Ravi Jagannathan, and David E. Runkle, 1993, On the relation between the expected value and the volatility of the nominal excess return on stocks, *Journal of Finance* 48, 1779{1801.
- Hall, Robert E., 1988, Intertemporal substitution in consumption, *Journal of Political Economy* 96, 339{57.
- Hamilton, James D., 1989, A new approach to the economic-analysis of nonstationary time-series and the business-cycle, *Econometrica* 57, 357{384.
- , 1994, *Time Series Analysis* (Princeton University Press).

- Hansen, Lars Peter, John C. Heaton, and Nan Li, 2008, Consumption strikes back? Measuring long-run risk, *Journal of Political Economy* 116, 260{302.
- Hodrick, R. J., 1992, Dividend yields and expected stock returns - alternative procedures for inference and measurement, *Review of Financial Studies* 5, 357{386.
- Jacobs, Kris, and Kevin Q. Wang, 2004, Idiosyncratic consumption risk and the cross section of asset returns, *Journal of Finance* 59, 2211{2252.
- Jagannathan, Ravi, and Yong Wang, 2007, Lazy investors, discretionary consumption, and the cross-section of stock returns, *Journal of Finance* 62, 1623{1661.
- Jagannathan, Ravi, and Zhenyu Wang, 1996, The conditional CAPM and the cross-section of expected returns, *Journal of Finance* 51, 3{53.
- Kandel, Shmuel, and Robert F. Stambaugh, 1990, Expectations and volatility of consumption and asset returns, *Review of Financial Studies* 3, 207{32.
- , 1991, Asset returns and intertemporal preferences, *Journal of Monetary Economics* 27, 39{71.
- Keim, Donald B., and Robert F. Stambaugh, 1986, Predicting returns in the stock and bond markets, *Journal of Financial Economics* 17, 357{390.
- Kim, Chang-Jin, and Charles R. Nelson, 1999, Has the us economy become more stable? a bayesian approach based on a markov-switching model of the business cycle, *Review of Economics and Statistics* 81, 608{616.
- Kreps, David M., and Evan L. Porteus, 1978, Temporal resolution of uncertainty and dynamic choice theory, *Econometrica* 46, 185{200.
- Lettau, Martin, and Sydney Ludvigson, 2001a, Consumption, aggregate wealth, and expected stock returns, *Journal of Finance* 56, 815{849.
- , 2001b, Resurrecting the (C)CAPM: A cross-sectional test when risk premia are time-varying, *Journal of Political Economy* 109, 1238{1287.
- , and Jessica Wachter, 2008, The declining equity premium: What role does macroeconomic risk play?, *Review of Financial Studies* 21, 1607{1652.
- Lewellen, Jonathan, 2004, Predicting returns with financial ratios, *Journal of Financial Economics* 74, 209{235.
- Lustig, Hanno, and Stijn Van Nieuwerburgh, 2005, Housing collateral, consumption insurance and risk premia: An empirical perspective, *Journal of Finance* 60, 1167{1219.
- Lustig, Hanno, Stijn Van Nieuwerburgh, and Adrien Verdelhan, 2008, The wealth-consumption ratio, Working Paper.
- Mankiw, N. Gregory, and Matthew D. Shapiro, 1986, Risk and return: Consumption beta versus market beta, *Review of Economics and Statistics* 68, 452{459.
- Newey, Whitney K., and Kenneth D. West, 1987, A simple, positive semidefinite, heteroskedasticity and autocorrelation consistent covariance-matrix, *Econometrica* 55, 703{708.
- Parker, Jonathan A., and Christian Julliard, 2005, Consumption risk and the cross section of expected returns, *Journal of Political Economy* 113, 185{222.
- Piazzesi, Monika, Martin Schneider, and Selale Tuzel, 2007, Housing, consumption, and asset pricing, *Journal of Financial Economics* 83, 531{569.

- Pindyck, Robert S., 1984, Risk, in ation, and the stock market, *The American Economic Review* 74, 335{351.
- Poterba, James M., and Lawrence H. Summers, 1986, The persistence of volatility and stock market uctuations, *The American Economic Review* 76, 1142{1151.
- Restoy, Fernando, 1991, Essays in intertemporal asset pricing, PhD dissertation, Harvard University.
- , and Philippe Weil, 2004, Approximate equilibrium asset prices, Working paper.
- Roze , Michael S., 1984, Dividend yields are equity risk premiums, *Journal of Portfolio Management* 11, 68{75.
- Santos, Tano, and Pietro Veronesi, 2006, Labor income and predictable stock returns, *Review of Financial Studies* 19, 1{44.
- Shanken, Jay, 1992, On the estimation of beta-pricing models, *Review of Financial Studies* 5, 1{33.
- Stambaugh, Robert F., 1999, Predictive regressions, *Journal of Financial Economics* 54, 375{421.
- Tedongap, Romeo, 2007, Consumption volatility and the cross-section of stock returns, Working Paper.
- Vassalou, Maria, 2003, News related to future gdp growth as a risk factor in equity returns, *Journal of Financial Economics* 68, 47{73.
- Veronesi, Pietro, 1999, Stock market overreaction to bad news in good times: A rational expectations equilibrium model, *Review of Financial Studies* 12, 975{1007.
- Vissing-Jorgensen, A., 2002, Limited asset market participation and the elasticity of intertemporal substitution, *Journal of Political Economy* 110, 825{853.
- , and O. P. Attanasio, 2003, Stock-market participation, intertemporal substitution, and risk-aversion, *American Economic Review* 93, 383{391.
- Weil, Philippe, 1989, The equity premium puzzle and the risk-free rate puzzle, *Journal of Monetary Economics* 24, 401{421.
- Whitelaw, Robert F., 2000, Stock market risk and return: An equilibrium approach, *Review of Financial Studies* 13, 521{47.
- Wilcox, David W., 1992, The construction of u.s. consumption data: Some facts and their implications for empirical work, *The American Economic Review* 82, 922{941.
- Yogo, Motohiro, 2006, A consumption-based explanation of expected stock returns, *Journal of Finance* 61, 539{580.

Figure 1. Bayesian Beliefs about the Mean and Volatility State

This figure displays the estimated Bayesian belief process for the state of the conditional mean (top figure) and conditional volatility (bottom figure) of consumption growth. The estimation procedure follows Hamilton (1994). We use quarterly per capita real non-durable plus service consumption.




Figure 2. Wealth-Consumption Ratio for a high EIS Agent

This figure shows the wealth-consumption ratio as a function of the perceived conditional first $\hat{\mu}_t$ (left graph) and second $\hat{\sigma}_t$ (right graph) moment of consumption growth for the benchmark calibration when the agent has a high EIS of 1.5. Further, relative risk aversion equals 30 and the quarterly rate of time preference is 0.995.




Figure 3. Wealth-Consumption Ratio for a low EIS Agent

This figure shows the wealth-consumption ratio as a function of the perceived conditional first $\hat{\mu}_t$ (left graph) and second $\hat{\sigma}_t$ (right graph) moment of consumption growth for the benchmark calibration when the agent has a low EIS of 0.5. Further, relative risk aversion equals 30 and the quarterly rate of time preference is 0.995.




Figure 4. Pricing Errors of the Consumption-Based Model

This figure depicts average quarterly excess returns of the 25 Fama-French portfolios against model predicted excess returns. In all graphs, consumption growth (Δc_t) is included as explanatory factor. In the top-right and bottom-left graph, we add changes in the perceived second ($\Delta \sigma_t$) and first ($\Delta \mu_t$) moments of consumption growth. In the bottom-right graph, we replace beliefs about the volatility state with the VR factor in the full model.

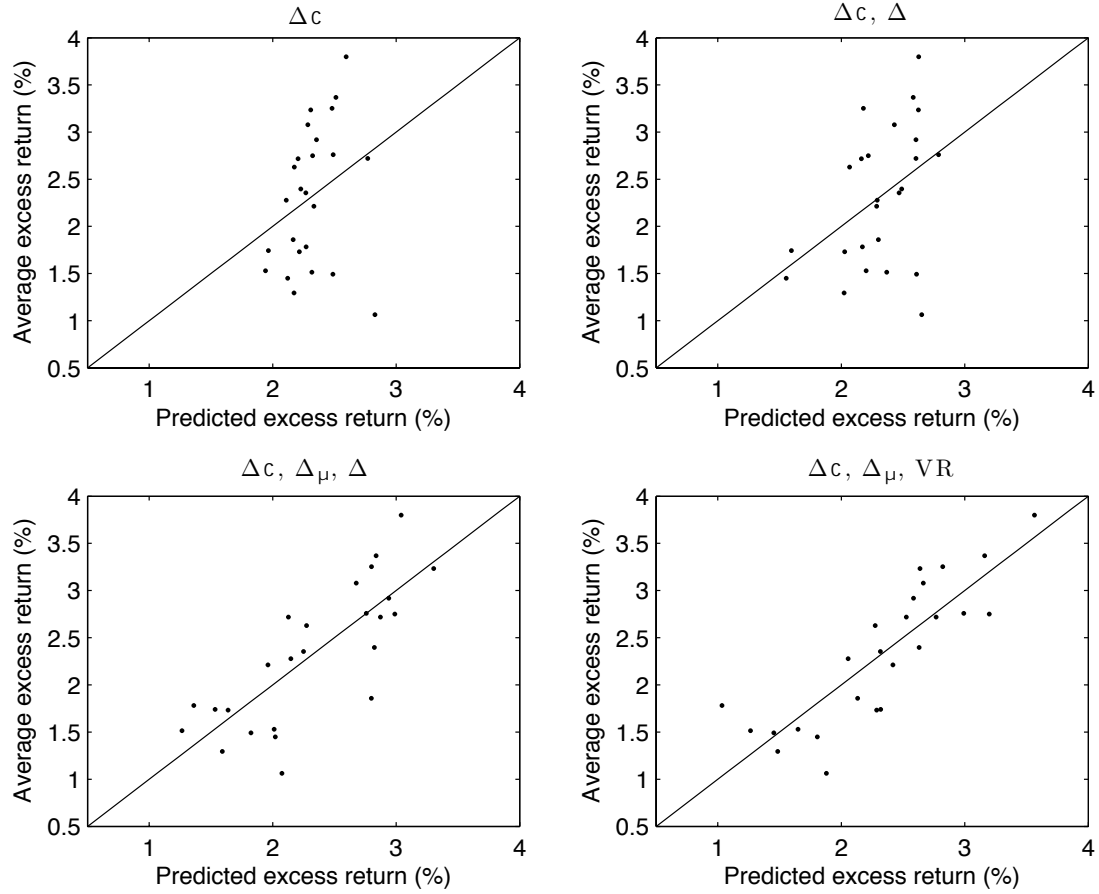


Figure 5. Pricing Errors of the Market-Based Model

This figure depicts average monthly excess returns of the 25 Fama-French portfolios against model predicted excess returns. The top-left graph represents the standard CAPM and the bottom-left graph the Fama-French three factor model. In the top-right graph, we add the VR factor to the CAPM and, in the bottom-right graph, we replace the HML factor with VR in the Fama-French model.

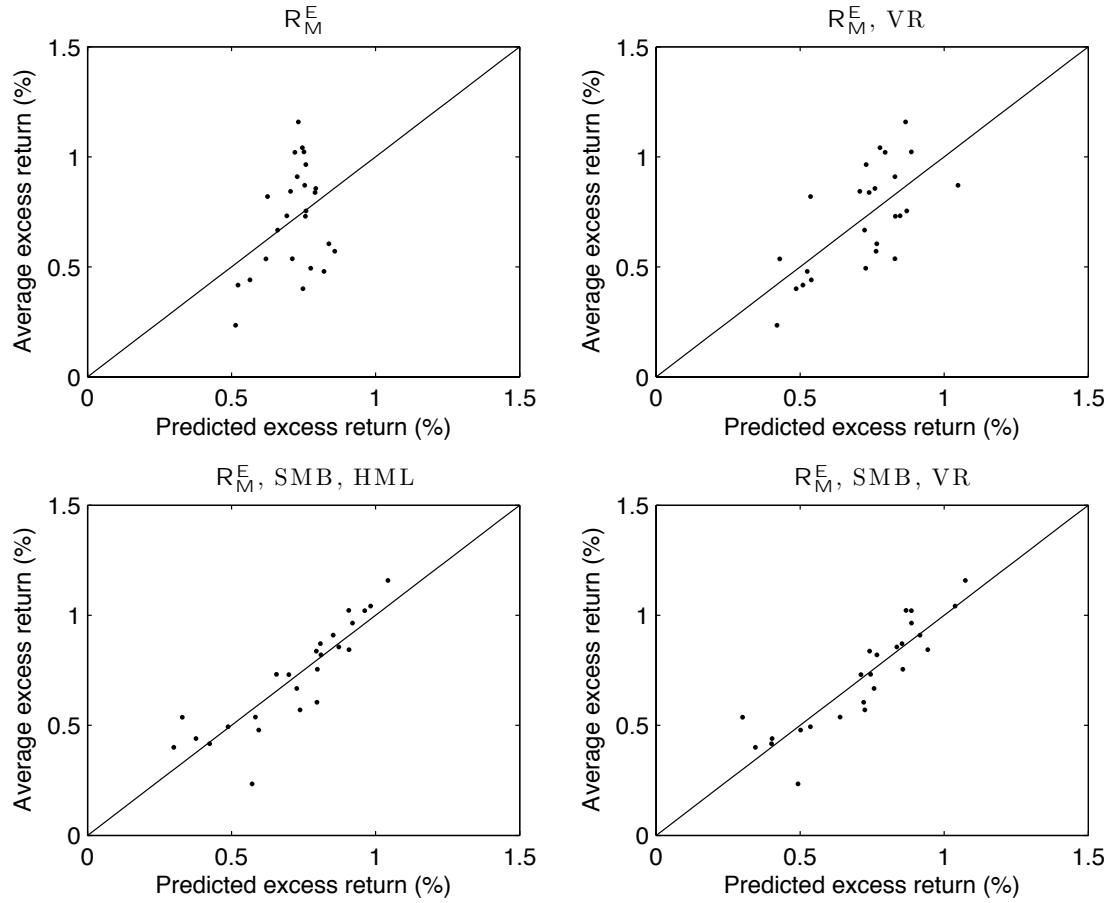


Table I
Markov Model of Consumption Growth

This table reports parameter estimates of the Markov model for log consumption growth

$$\Delta c_{t+1} = \mu_t + \sigma_t \epsilon_{t+1} \quad \epsilon_t \sim \mathcal{N}(0; 1)$$

where $\mu_t \in \{l; h\}$ and $\sigma_t \in \{l; h\}$ follow independent Markov processes with transition matrices P^μ and P^σ respectively.

$$P^\mu = \begin{bmatrix} p_{ll}^\mu & 1 - p_{ll}^\mu \\ 1 - p_{hh}^\mu & p_{hh}^\mu \end{bmatrix} \quad P^\sigma = \begin{bmatrix} p_{ll}^\sigma & 1 - p_{ll}^\sigma \\ 1 - p_{hh}^\sigma & p_{hh}^\sigma \end{bmatrix}$$

Panel C shows the correlation between the filtered beliefs for each state.

Panel A: Parameter Estimates			
μ	σ	μ	σ
0.375	0.748	0.211	0.463
(0.034)	(0.050)	(0.023)	(0.042)
Panel B: Marginal Transition Probabilities			
p_{μ}^{ll}	p_{μ}^{hh}	p_{σ}^{ll}	p_{σ}^{hh}
0.935	0.911	0.977	0.978
(0.037)	(0.051)	(0.028)	(0.024)
Panel C: Correlation of Beliefs			
	b_μ	b_σ	
b_μ	1		
b_σ	0.149	1	

Table II
Fama-MacBeth Regressions

This table reports cross-sectional regressions of monthly returns on estimated risk loadings and characteristics. Time-varying risk loadings are obtained from 5-year rolling time-series regressions of individual excess returns on the market excess return, log consumption growth, and changes in the perceived conditional mean and volatility of consumption growth using quarterly data. In the cross-section, we regress monthly future returns onto the loadings on market excess return ($\hat{\mu}_{M,t}^i$), log consumption growth ($\hat{c}_{c,t}^i$), changes in the perceived conditional mean ($\hat{\mu}_{\mu,t}^i$) and volatility ($\hat{\sigma}_{\sigma,t}^i$) of consumption growth as well as market capitalization (ME_t^i) and book-to-market ratio (BM_t^i). Both characteristics are measured in December which is 7 to 18 months prior the test month. All explanatory variables are normalized so they are centered around zero with unit variance. Reported are time-series averages of the second stage coefficients with associated t-statistics. The sample period is from January 1964 to December 2007. t-statistics are in parenthesis.

Fama-MacBeth Regressions						
MODEL	$\hat{\mu}_{M,t}^i$	$\hat{c}_{c,t}^i$	$\hat{\mu}_{\mu,t}^i$	$\hat{\sigma}_{\sigma,t}^i$	ME_t^i	BM_t^i
I	0.026 (0.33)					
II		0.043 (0.85)				
III			0.000 (0.00)			
IV				-0.134 (-3.23)		
V		0.159 (1.98)	0.033 (0.57)	-0.151 (-2.45)		
VI		0.178 (2.36)	0.044 (0.82)	-0.113 (-2.10)	-0.057 (-1.74)	0.490 (2.54)
VII	0.034 (0.40)		-0.067 (-1.32)	-0.109 (-2.50)		
VIII	0.021 (0.26)		-0.068 (-1.42)	-0.089 (-2.25)	-0.067 (-2.32)	0.543 (2.81)
IX	0.042 (0.49)	0.142 (1.81)	0.018 (0.30)	-0.119 (-2.17)		
X	0.039 (0.46)	0.160 (2.12)	0.028 (0.48)	-0.086 (-1.73)	-0.066 (-2.34)	0.503 (2.61)

Table III
Portfolios formed on Risk Exposure

Each quarter stocks are assigned into quintiles based on loadings from time series regressions (18) of individual excess returns on the market excess return ($\hat{M}_{t,t}^i$), log consumption growth ($\hat{c}_{t,t}^i$), and changes in the perceived first ($\hat{\mu}_{t,t}^i$) and second moment ($\hat{\sigma}_{t,t}^i$) of consumption growth. This table reports average equally-weighted (EW) and value-weighted (VW) returns of these portfolios and their associated t-statistics.

Panel A: Univariate sorts based on $\hat{M}_{t,t}^i$						
$\hat{M}_{t,t}^i$	low		med		high	high - low
EW	1.20 (6.78)	1.26 (6.50)	1.31 (5.90)	1.33 (5.09)	1.28 (3.89)	0.09 (0.40)
VW	0.98 (6.15)	0.94 (5.38)	0.90 (4.67)	0.90 (3.90)	0.91 (3.14)	-0.07 (-0.33)
Panel B: Univariate sorts based on $\hat{c}_{t,t}^i$						
$\hat{c}_{t,t}^i$	low		med		high	high - low
EW	1.21 (4.56)	1.23 (5.78)	1.25 (6.02)	1.31 (5.88)	1.37 (4.85)	0.16 (1.28)
VW	0.97 (4.54)	0.97 (5.40)	0.90 (5.08)	0.92 (4.77)	1.05 (4.20)	0.08 (0.46)
Panel C: Univariate sorts based on $\hat{\mu}_{t,t}^i$						
$\hat{\mu}_{t,t}^i$	low		med		high	high - low
EW	1.34 (5.02)	1.23 (5.69)	1.26 (6.17)	1.22 (5.61)	1.31 (4.73)	-0.02 (-0.21)
VW	1.04 (4.54)	0.95 (5.04)	0.96 (5.60)	0.87 (4.73)	1.04 (4.52)	0.00 (-0.01)
Panel D: Univariate sorts based on $\hat{\sigma}_{t,t}^i$						
$\hat{\sigma}_{t,t}^i$	low		med		high	high - low
EW	1.46 (5.26)	1.33 (6.04)	1.27 (6.14)	1.17 (5.48)	1.15 (4.41)	-0.32 (-2.92)
VW	1.18 (5.33)	0.99 (5.36)	0.96 (5.30)	0.99 (5.18)	0.75 (3.27)	-0.43 (-3.02)

Table IV
Characteristics of Volatility Risk Exposure Portfolios

This table reports various characteristics and risk measures for the quintile portfolios based on consumption volatility loadings (as in Table III, Panel D). Panel A shows average returns and consumption volatility betas. Panel B reports the average value and mean decile rank for size and book-to-market characteristics of each portfolio.

Panel A: Univariate sorts based on $\hat{\sigma}_{\sigma,t}^i$			
low	med	high	high - low

Table V
Portfolios Formed on Consumption Volatility Risk and Market Capitalization

This table reports average returns of equally-weighted (Panel A) and value-weighted (Panel B) portfolios formed independently on consumption volatility exposure (low 30%, medium 40%, and high 30%) and market capitalization (small 50% and big 50%).

Panel A: Equally-Weighted Returns				
	Low	Med	High	H - L
Small	1.57 (5.35)	1.50 (6.05)	1.37 (4.93)	-0.20 (-2.12)
Big	1.24 (5.33)	1.15 (5.79)	0.99 (4.20)	-0.24 (-2.75)
S - B	0.33 (2.09)	0.35 (2.78)	0.37 (2.51)	
Panel B: Value-Weighted Returns				
	Low	Med	High	H - L
Small	1.47 (5.45)	1.43 (6.02)	1.27 (4.74)	-0.20 (-1.85)
Big	1.21 (5.57)	1.08 (5.73)	0.96 (4.24)	-0.25 (-2.73)
S - B	0.26 (1.95)	0.34 (3.09)	0.32 (2.35)	

Table VI
Portfolios Formed on Consumption Volatility Risk and Book-to-Market Ratio

This table reports average returns of equally-weighted (Panel A) and value-weighted (Panel B) portfolios formed independently on consumption volatility exposure (low 30%, medium 40%, and high 30%) and the ratio of book equity to market equity (low 50% and high 50%).

Panel A: Equally-Weighted Returns				
	Low	Med	High	H - L
Low BM	1.19 (4.56)	1.04 (4.76)	0.84 (3.34)	-0.35 (-3.36)
High BM	1.52 (6.13)	1.35 (6.73)	1.31 (5.47)	-0.22 (-2.55)
H - L	0.34 (3.63)	0.31 (3.77)	0.47 (4.58)	
Panel B: Value-Weighted Returns				
	Low	Med	High	H - L
Low BM	1.03 (4.93)	0.91 (4.80)	0.80 (3.66)	-0.24 (-1.85)
High BM	1.29 (6.29)	1.06 (6.19)	0.99 (4.65)	-0.30 (-2.15)
H - L	0.26 (2.17)	0.15 (1.44)	0.20 (1.46)	

Table VII
Volatility Risk Pricing

This table reports second stage regressions to estimate market prices to risk. Δc_t denotes log consumption growth, $\Delta_{-1} c_t$ and $\Delta_{-2} c_t$ are changes in filtered beliefs about the first and second moment of consumption growth. The VR factor is the return of holding a long position in the value-weighted quintile of stocks with high volatility risk ($\hat{\sigma}_{\sigma,t}^i$) and a short position in low volatility risk, as reported in Panel D of Table III. Test assets are the value weighted 25 Fama-French portfolios constructed from independent quintile sorts based on market capitalization and book-to-market ratio. The quarterly time-series starts in 1964 and ends 2007. The t-statistics are corrected for estimation error in the first stage as proposed by Shanken (1992) and are Newey and West (1987) adjusted to account for heteroskedasticity and autocorrelation. For each regression, the last two columns report R^2 (and adjusted R^2) and the regression J-statistic (χ^2) with the associated p-value.

	CONST.	Δc_t	$\Delta_{-1} c_t$	$\Delta_{-2} c_t$	VR	R^2	χ^2
	(t-stat)	(t-stat)	(t-stat)	(t-stat)	(t-stat)	(adj. R^2)	(p-val)
I	1.61 (2.38)	0.17 (0.86)				8.39 (4.41)	81.83 (0.00)
II	2.05 (2.01)	0.46 (1.89)	-0.01 (-0.44)			44.10 (39.02)	39.62 (1.69)
III	2.12 (1.94)	-0.03 (-0.11)		-0.04 (-1.94)		17.61 (10.12)	28.98 (18.10)
IV	2.85 (1.48)	0.22 (0.65)	-0.05 (-0.90)	-0.07 (-1.55)		62.18 (56.78)	11.86 (96.02)
V	1.10 (1.02)	0.31 (1.46)	0.02 (0.69)		-6.01 (-2.46)	72.31 (68.36)	34.03 (4.88)
VI	1.21 (1.31)	0.34 (1.32)			-5.70 (-2.24)	71.98 (69.44)	37.80 (2.68)

Table VIII
Volatility Risk Factor

This table provides descriptive statistics of the volatility risk (VR) portfolio. The VR portfolio is the return of holding a long position in the value-weighted quintile of stocks with high volatility risk ($\hat{\sigma}_{\sigma,t}^i$) and a short position in low volatility risk, as reported in Panel D of Table III. The time-series for the VR portfolio covers January 1964 until December 2007. Panel A shows the mean, standard deviation and the Sharpe Ratio of the VR portfolio as well as the three Fama-French portfolios. Panel B presents the correlation matrix of the factor returns, and Panel C reports parameter estimates from time-series regressions of the VR portfolio on the market and the three Fama-French factors.

Panel A: Summary Statistics				
	VR	MKTMRF	SMB	HML
N	528	528	528	528
MEAN	-0.43	0.46	0.25	0.43
STD	3.25	4.37	3.22	2.91
SR	-0.13	0.11	0.08	0.15
Panel B: Correlations				
	VR	MKTMRF	SMB	HML
VR	1			
MKTMRF	0.11	1		
SMB	0.08	0.30	1	
HML	-0.23	-0.41	-0.27	1
Panel C: Time-Series Regressions				
	$\hat{\alpha}$	$\hat{\alpha}_{MKTMRF}$	$\hat{\alpha}_{HML}$	$\hat{\alpha}_{SMB}$
I	-0.43 (-3.02)			
II	-0.47 (-3.29)	0.08 (2.56)		
III	-0.33 (-2.32)	0.01 (0.29)	-0.25 (-4.73)	0.02 (0.41)

Table IX
Factor Exposures of the 25 Fama-French Portfolios

This table reports factor loadings of the 25 Fama-French portfolios with the market return and volatility risk (VR) portfolio. The VR portfolio is the return of holding a long position in the value-weighted quintile of stocks with high volatility risk ($\hat{\sigma}_{i,t}^h$) and a short position in low volatility risk, as reported in Panel D of Table III. Test assets are the value-weighted 25 Fama-French portfolios constructed from independent quintile sorts based on market capitalization (S1: small, S5: big) and book-to-market ratio (BM1: low, BM5: high). The time series starts in January 1964 and ends in December 2007.

Panel A: Loading on the Market Factor						
	BM1	BM2	BM3	BM4	BM5	Average
S1	1.45 (29.17)	1.24 (28.57)	1.09 (30.45)	1.01 (29.51)	1.04 (27.77)	1.17
S2	1.44 (38.58)	1.18 (39.11)	1.05 (37.90)	0.99 (35.53)	1.06 (32.25)	1.14
S3	1.36 (43.52)	1.12 (48.90)	0.99 (42.72)	0.92 (37.74)	1.01 (33.24)	1.08
S4	1.25 (53.66)	1.08 (56.88)	0.99 (46.68)	0.93 (40.40)	1.01 (34.84)	1.05
S5	1.00 (57.97)	0.96 (57.31)	0.86 (42.01)	0.79 (33.95)	0.83 (27.53)	0.89
Average	1.30	1.12	1.00	0.93	0.99	
Panel B: Loading on the Volatility Risk Factor						
	BM1	BM2	BM3	BM4	BM5	Average
S1	0.14 (2.14)	0.04 (0.74)	-0.06 (-1.18)	-0.10 (-2.24)	-0.13 (-2.50)	-0.02
S2	0.11 (2.17)	-0.04 (-0.91)	-0.11 (-2.97)	-0.09 (-2.43)	-0.09 (-2.11)	-0.04
S3	0.08 (1.84)	-0.10 (-3.13)	-0.14 (-4.47)	-0.12 (-3.62)	-0.14 (-3.48)	-0.08
S4	0.09 (2.83)	-0.10 (-3.87)	-0.12 (-4.31)	-0.11 (-3.57)	-0.20 (-5.16)	-0.09
S5	0.00 (0.00)	-0.10 (-4.40)	-0.05 (-1.86)	-0.16 (-4.98)	-0.15 (-3.59)	-0.09
Average	0.08	-0.06	-0.10	-0.12	-0.14	

Table X
Volatility Risk Pricing Factor

This table reports statistics from two-pass regressions to document the pricing implications of the volatility risk (VR) model. The VR factor is the return of holding a long position in the value-weighted quintile of stocks with high volatility risk ($\hat{\sigma}_{\sigma,t}^i$) and a short position in low volatility risk, as reported in Panel D of Table III. Test assets are the value weighted 25 Fama-French portfolios constructed from independent quintile sorts based on market capitalization and book-to-market ratio. The time-series covers January 1964 until December 2007. The t-statistics are corrected for estimation error in the first stage as proposed by Shanken (1992) and are Newey and West (1987) adjusted with 6 lags to account for heteroskedasticity and autocorrelation. For each regression, the last two columns report R^2 (and adjusted R^2) and the regression J-statistic (J^2) with the associated p-value (in %).

	CONST (t-stat)	R_M^E (t-stat)	SMB (t-stat)	HML (t-stat)	VR (t-stat)	R^2 (adj. R^2)	J^2 (p-val)
I	1.25 (3.23)	-0.50 (-1.19)				14.08 (10.35)	67.77 (0.00)
II	0.83 (3.81)				-1.21 (-1.72)	32.86 (29.94)	62.22 (0.00)
III	-0.30 (-0.49)	0.77 (1.18)			-2.84 (-2.76)	46.27 (41.38)	40.15 (1.48)
IV	1.22 (3.81)	-0.72 (-1.94)	0.20 (1.34)	0.46 (3.03)		77.38 (74.15)	53.10 (0.02)
V	1.11 (3.27)	-0.67 (-1.73)	0.20 (1.33)		-1.80 (-2.51)	79.92 (77.05)	41.43 (0.73)
VI	1.02 (3.01)	-0.57 (-1.48)	0.20 (1.37)	0.44 (2.93)	-1.45 (-2.26)	80.30 (76.36)	44.52 (0.20)
VII	-0.26 (-0.48)	0.80 (1.36)		0.39 (2.54)		48.37 (43.67)	63.28 (0.00)
VIII	-0.46 (-0.79)	0.95 (1.51)		0.37 (2.43)	-1.46 (-2.27)	51.50 (44.57)	54.81 (0.01)

Table XI
Market Predictability in the Time-Series

This table reports time-series regressions of excess returns on lagged predictor variables. The market return is the value-weighted CRSP index less the 90 day T-bill rate. Predictor variables are the lagged market return ($R_{M,t-1}$), the 90 day T-bill less its 12 months moving average (RelTbill), the difference between yields of long and short government bonds (Term), the yield difference between Baa rated and Aaa rated corporate bonds (Default), the dividend yield (DY), the consumption-wealth ratio of Lettau and Ludvigson (2001a) (cay), and changes in beliefs about the first (Δ_t) and second moments (Δ_t) of consumption growth. The sample period includes the first quarter of 1957 until the fourth quarter of 2006.

	$R_{M,t-1}$	RelTbill	Term	Default	DY	cay	Δ_t	Δ_t	R^2 (%)
I	-0.02 (-0.28)	-0.01 (-1.57)	0.01 (0.83)	0.00 (-0.16)	0.93 (1.39)				4.58
II	0.05 (0.66)					1.51 (3.50)			5.97
III	0.01 (0.18)	-0.01 (-1.64)	0.00 (-0.43)	0.01 (0.66)	0.05 (0.07)	1.44 (2.87)			8.50
IV	0.03 (0.44)						0.17 (0.02)		0.11
V	0.04 (0.55)							59.94 (3.09)	4.71
VI	-0.01 (-0.20)	-0.01 (-1.44)	0.01 (1.12)	-0.01 (-0.41)	1.03 (1.57)			60.02 (3.11)	9.13
VII	0.05 (0.76)					1.50 (3.55)		59.35 (3.14)	10.48
VIII	0.02 (0.23)	-0.01 (-1.50)	0.00 (-0.12)	0.01 (0.39)	0.20 (0.28)	1.36 (2.75)		56.91 (2.99)	12.57