

Consumption Volatility Risk*

Oliver Boguth
Sauder School of Business
University of British Columbia

Lars-Alexander Kuehn
Tepper School of Business
Carnegie Mellon University

November 19, 2009

ABSTRACT

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. In a model with recursive preferences and unobservable conditional mean and volatility of consumption growth, prior beliefs about conditional moments of consumption growth affect excess returns. Empirically, we find that perceived consumption volatility is a priced source of risk and exposure to it negatively predicts future returns in the cross-section. In the time-series, beliefs about the volatility state strongly forecast aggregate quarterly excess returns.

JEL Classification: G12, G17, E44.

Keywords: Asset pricing, consumption volatility, cross-section of returns, volatility risk, recursive preferences.

*We thank Murray Carlson, Alexander David, Wayne Ferson, Adlai Fisher, Lorenzo Garlappi, Burton Hollifield, Lars Lochstoer (WFA discussant), Monika Piazzesi, Lukas Schmid (EFA discussant), Amir Yaron, Motohiro Yogo, Lu Zhang and seminar participants at Berkeley, UBC and Wharton as well as conference participants at the 2009 North American Summer Meetings of the Econometric Society, 2009 WFA meeting, 2009 CEPR Gerzensee Summer Symposium on Financial Markets, 2009 EFA meeting and 2009 NFA meeting for valuable comments. We also thank Kenneth French, Amit Goyal and Martin Lettau for providing data on their website. Contact information, Boguth: 2053 Main Mall, Vancouver BC, Canada V6T 1Z2, oliver.boguth@sauder.ubc.ca; Kuehn: 5000 Forbes Avenue, Pittsburgh, PA 15213, kuehn@cmu.edu.

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), the impact of consumption growth volatility on asset prices has received less attention in the empirical literature.¹ This is surprising since it is well known that the volatility of macroeconomic quantities, such as consumption and output, varies over time.² The goal of this paper is to analyze the pricing implications of consumption growth volatility in the cross-section and time-series of stock returns.

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 has no separate innovations relative to the process for consumption growth. 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. Last, aggregate consumption is measured with error thereby making statistical inference more difficult (Breedon, Gibbons, and Litzenberger (1989) and Wilcox (1992)).

Our model follows the work of Bansal and Yaron (2004) and Kandel and Stambaugh (1991) and uses the same building blocks as Lettau, Ludvigson, and Wachter (2008). 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. An important implication of recursive preferences is that the agent cares not only about shocks to current consumption growth but also about changes to the conditional distribution of future consumption growth. In our model, these changes are driven by persistent states of the

¹A notable exception is Bansal, Kiku, and Yaron (2007). Following their model, they estimate the conditional first and second moments of consumption growth as affine functions of financial data. Tedongap (2007) uses a GARCH process for consumption volatility. Other recent contributions testing the C-CAPM using realized consumption growth 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).

²For instance, see Cecchetti and Mark (1990), Kandel and Stambaugh (1990), Bonomo and Garcia (1994), Kim and Nelson (1999), or Whitelaw (2000).

Markov chains for the first and second moments of consumption growth. We further assume that the state of the economy is unobservable and the agent uses Bayesian updating to form beliefs about the state, similar to David (1997) and Veronesi (1999). As a result, the agent's estimates of the conditional first and second moments of consumption growth are priced.

The model has the following implication for the cross-section of returns. When the elasticity of intertemporal substitution (EIS) is greater than the inverse of the coefficient of relative risk aversion (RRA), the agent prefers intertemporal risk due to unobservable Markov states to be resolved sooner rather than later. 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 learn that future consumption is (not) very volatile. This asset serves as insurance against uncertain times and thus has a lower required return. Consequently, the agent demands a positive market price of risk for shocks to expected consumption growth and a negative one for shocks to the conditional volatility of consumption growth. To provide convincing empirical evidence, we test these implications in two ways. First, we study the relation between risk loadings and future returns at the firm level.³ Second, we estimate the market price of risk directly using portfolios.

Following Hamilton (1989), we estimate a Markov chain for the first and second moments of consumption growth. Bayesian updating provides beliefs about the states for mean and volatility. To obtain time-varying risk loadings with respect to innovations in the perceived conditional first and second moments of consumption growth, we run rolling quarterly time-series regressions of individual stock returns on consumption growth as well as innovations in beliefs for mean and volatility. Sorting stocks into portfolios based on these risk loadings, we find that loadings on innovations in the perceived expected consumption growth do not help to explain future returns. Loadings on consumption growth volatility, however, significantly

³Ang, Liu, and Schwarz (2008) emphasize the use of firm level data to estimate market prices of risk because firm level data display more dispersion in betas. As a result, the estimation is more efficient.

negatively forecast cross-sectional differences in returns. A consumption volatility risk factor (CVR), 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 average book-to-market ratios.

The negative relation between consumption volatility risk loadings and future returns at the firm level suggests a negative price of consumption volatility risk. In order to provide direct evidence, we perform two stage regressions of excess returns on log consumption growth, changes in the perceived mean and volatility of consumption and the CVR factor. Importantly, the coefficients on the innovation in the perceived consumption volatility and CVR are negative implying that the representative agent has an EIS greater than the inverse of RRA. A crucial assumption in the long-run risk framework of Bansal and Yaron (2004) is that the agent prefers intertemporal risk to resolve sooner than later. Our findings strongly support this assumption.

We also augment the market CAPM and Fama-French 3-factor model with the CVR factor. In particular, CVR shows up strongly and significant in addition to the market and the three Fama and French (1993) factors. When the CVR factor is added to specifications that already contain the value factor HML, average absolute pricing errors decline only marginally. At the same time, replacing HML with CVR does not result in larger average pricing errors. We thus conclude that HML and CVR have similar pricing implications. But in contrast to HML, the volatility risk factor has a clear economic interpretation.

Another implication of our model is the predictability of the aggregate 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 an 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. A one standard deviation increase of the perceived consumption volatility results in a 1.4% rise of the quarterly equity premium, similar to the predictive power of the wealth-consumption 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 wealth-consumption ratio. Thus, one might expect that direct measures of the wealth-consumption ratio, such as *cay*, comprise all relevant 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.

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. The positive relation between consumption volatility and the equity premium provides evidence for an EIS greater than one.

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.⁴ 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

⁴One 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).

volatility feedback effect. Similarly, French, Schwert, and Stambaugh (1987), 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 a short- and long-run component and show how each of the two components affects the cross-section of asset prices.

All of the above papers use some measure of stock market volatility. Motivated by the long-run risk model of Bansal and Yaron (2004), several important papers study the relation between consumption volatility and prices. Notably, Bansal, Khatchatrian, and Yaron (2005) find that the conditional consumption volatility predicts aggregate valuation ratios. Bansal, Kiku, and Yaron (2007) estimate the long-run risk model using the cross-section of returns. Following their theory, they estimate consumption volatility as an affine function of the observable aggregate price-dividend ratio and short-term interest rate. Their Table IV indicates that consumption volatility plays a minor role in explaining the size and value spread relative to shocks to expected consumption growth. In contrast, we filter consumption volatility directly from consumption data without the use of financial data. We find that consumption volatility is a dominant contributor to risk premia in the cross-section.⁵

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 premia. We extend their work by studying the cross-section and time-series of returns. Closely related is also Calvet and Fisher (2007) who study the asset pricing implications of multi-fractal Markov switching in a recursive preference model at the aggregate level.

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

⁵Jacobs and Wang (2004) and Balduzzi and Yao (2007) 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.

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. 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.

Drechsler and Yaron (2008) extend the long-run risk model to include jumps in consumption growth and volatility. Their model generates a variance premium and return predictability which are 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.⁶

The remainder of the paper is organized as follows: In Section I, 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 II-IV. In Section II, we test whether loadings on consumption growth and its conditional moments forecast returns in the cross-section. We form portfolios and run Fama-MacBeth regressions based on consumption volatility loadings. In Section III, we test whether consumption growth and its conditional moments as well as the CVR factor are priced risk factors. Section IV contains time-series predictability tests and Section V concludes. The appendix contains derivations and additional results.

⁶Other 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).

I. 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. In our model, future consumption growth is influenced by time-variation in its conditional mean and volatility and the agent's beliefs about the aggregate state enter the pricing kernel through the wealth-consumption ratio.

A. Consumption

We assume that the conditional first and second moments of consumption growth follow a Markov chain. Specifically, log consumption growth, Δc_{t+1} , follows

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

where μ_t denotes its conditional expectation and σ_t its conditional standard deviation. For tractability in the empirical estimation, we assume two states for the mean and two for the volatility which are denoted by $\mu_t \in \{\mu_l, \mu_h\}$ and $\sigma_t \in \{\sigma_l, \sigma_h\}$. The conditional first and second moments of consumption growth follow Markov chains with transition matrices 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 model parsimonious, we impose that mean and volatility states switch independently. Thus, the joint transition matrix is the product of the marginal transition probabilities for mean and volatility states and the 16-element matrix can be fully characterized by 4 parameters. Importantly, the assumption of independent switching probabilities does not imply that the beliefs about mean and volatility states are independent. Since we assume two drift and two volatility states, there are four states in total, $\{(\mu_l, \sigma_l), (\mu_l, \sigma_h), (\mu_h, \sigma_l), (\mu_h, \sigma_h)\}$, denoted by $s_t \in \{1, \dots, 4\}$. Our specification follows Kandel and Stambaugh (1991), 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 as in Lettau, Ludvigson, and Wachter (2008). This assumption ensures that the empirical exercise is in line with the model. The inference 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 . We denote by $\xi_{t+1|t}$ the date- t prior belief vector about tomorrow's states

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

where

$$\eta_t = \begin{bmatrix} f(\Delta c_t | \mu_{t-1} = \mu_l, \sigma_{t-1} = \sigma_l, Y_{t-1}) \\ f(\Delta c_t | \mu_{t-1} = \mu_l, \sigma_{t-1} = \sigma_h, Y_{t-1}) \\ f(\Delta c_t | \mu_{t-1} = \mu_h, \sigma_{t-1} = \sigma_l, Y_{t-1}) \\ f(\Delta c_t | \mu_{t-1} = \mu_h, \sigma_{t-1} = \sigma_h, Y_{t-1}) \end{bmatrix}$$

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

B. Recursive Utility

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

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

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

Epstein-Zin preferences provide a separation between the EIS and RRA. These two concepts are inversely related when the agent has power utility. Intuitively, the EIS measures the agent's willingness to postpone consumption over time, a notion well-defined under certainty. Relative risk aversion measures the agent's aversion to atemporal risk across states.

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

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

where $\theta = \frac{1-\gamma}{1-1/\psi}$ and $PC_t = P_t/C_t$ denotes the wealth-consumption ratio. For the empirical exercise, it is useful to study the log-linearized pricing kernel. Intuitively, the first (stochastic) term in the pricing kernel is consumption growth, C_{t+1}/C_t , and the second one the growth rate of the wealth-consumption ratio, PC_{t+1}/PC_t . Hence, a log-linear approximation of the pricing kernel implicit in (5) is given by

$$m_{t+1} \approx k - \gamma \Delta c_{t+1} - (1 - \theta) \Delta pc_{t+1} \quad (6)$$

where small letters are logs of capital letters and Δ denotes first differences.⁷ The log pricing kernel (6) 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

$$\mathbb{E}_t[R_{i,t+1}^e] \approx \gamma \text{Cov}_t(R_{i,t+1}, \Delta c_{t+1}) + (1 - \theta) \text{Cov}_t(R_{i,t+1}, \Delta pc_{t+1}) \quad (7)$$

In an endowment model which is solely driven by i.i.d. 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 Markov state implies that the agent's prior probabilities characterize the state of the economy. Consequently, the wealth-consumption ratio is a function of the agent's beliefs, i.e., $PC_t = PC(\xi_{t+1|t})$. In order to study how the wealth-consumption ratio changes with beliefs about the state, we further define the prior belief that the mean state is high tomorrow by

$$b_{\mu,t} = P(\mu_{t+1} = \mu_h | \mathcal{F}_t) = \xi_{t+1|t}(3) + \xi_{t+1|t}(4) \quad (8)$$

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

$$b_{\sigma,t} = P(\sigma_{t+1} = \sigma_h | \mathcal{F}_t) = \xi_{t+1|t}(2) + \xi_{t+1|t}(4) \quad (9)$$

⁷More precisely, a log-linear approximation is given by

$$m_{t+1} \approx (\theta \ln \beta - (1 - \theta)k_0) - \gamma \Delta c_{t+1} - (1 - \theta)(pc_{t+1} - k_1 pc_t)$$

where $pc_t = \ln(P_t/C_t)$ denotes the log wealth-consumption ratio and k_0, k_1 are constants. The value of k_1 is given by $k_1 = \overline{PC}/(\overline{PC} - 1) > 1$, where \overline{PC} is the mean wealth-consumption ratio. Lustig, Van Nieuwerburgh, and Verdelhan (2008) estimate the unconditional quarterly wealth-consumption ratio to be close to 351 implying that $k_1 = 1.003$. Consequently, the log pricing kernel can be closely approximated by the log growth rate of the wealth-consumption ratio.

conditional on the current information set \mathcal{F}_t . The univariate effects of changing beliefs about the volatility (mean) state while holding the smean (volatility) state constant can locally be approximated. We show in the appendix that, given a constant volatility, changes of the log wealth-consumption ratio are

$$\Delta pc_{t+1} \approx \Delta b_{\mu,t+1} \left(\frac{1}{\theta} \frac{PC_{\mu=\mu_h,\sigma}^\theta - PC_{\mu=\mu_l,\sigma}^\theta}{b_{\mu,t} PC_{\mu=\mu_h,\sigma}^\theta + (1 - b_{\mu,t}) PC_{\mu=\mu_l,\sigma}^\theta} \right) \quad (10)$$

where $PC_{\mu,\sigma}$ denotes the wealth-consumption ratio when expected consumption growth is μ and the consumption volatility is σ . Analogously, given a constant mean, changes of the log wealth-consumption ratio are

$$\Delta pc_{t+1} \approx \Delta b_{\sigma,t+1} \left(\frac{1}{\theta} \frac{PC_{\mu,\sigma=\sigma_h}^\theta - PC_{\mu,\sigma=\sigma_l}^\theta}{b_{\sigma,t} PC_{\mu,\sigma=\sigma_h}^\theta + (1 - b_{\sigma,t}) PC_{\mu,\sigma=\sigma_l}^\theta} \right) \quad (11)$$

Equations (10) and (11) illustrate that changes in the log wealth-consumption ratio are 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 (7).

C. Estimation

To estimate the model, we obtain data on quarterly per capita real consumption expenditures from the Bureau of Economic Analysis as the sum of nondurables and services. The data is seasonally adjusted using the X-12-ARIMA filter. Ferson and Harvey (1992) analyze the impact of using filtered consumption data on asset pricing tests. In the appendix, we provide evidence that Markov states in the raw data survive the filter and can be identified from seasonally adjusted data.⁸ In accordance with the observation that 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 2008. The choice of 1955 provides sufficient consumption observations before the beginning of 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 large in the

⁸The BEA stopped providing seasonally unadjusted quarterly data in 2005.

high state relative to the low state ($\mu_l = 0.37\%$, $\mu_h = 0.78\%$ quarterly). State conditional consumption volatilities are $\sigma_l = 0.21\%$ and $\sigma_h = 0.48\%$. The probability of remaining in a given regime for the mean is 0.93 in the low state and 0.89 in the high state. The volatility regimes are somewhat more persistent, with probabilities of 0.95 and 0.96, 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). While there are small differences in the consumption measure, including the year 2008 in our analysis greatly reduces the persistence of the volatility regimes.

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. This constraint, however, is not a significant restriction of consumption data. We also estimate a model that allows the Markov chains for mean and volatility of consumption growth to be dependent. This unrestricted model has 8 additional parameters in the joint transition matrix. Yet the likelihood improves only marginally relative to the restricted model with independent Markov chains. A likelihood ratio test cannot be rejected at any significance level.⁹

Figure 1 shows the filtered beliefs for the regimes. Panel A depicts the belief dynamics for mean consumption growth $b_{\mu,t}$ and Panel B for the standard deviation $b_{\sigma,t}$. These graphs visually confirm that the mean regimes are less persistent than the volatility states. In particular, the parameter estimates for the Markov chain imply that mean states last for 3.1 years whereas volatility states last for 5.7 years on average. Further, a decline in consumption volatility from the 1990s onwards, as pointed out by Kim and Nelson (1999), is easily observable. The 2008 recession demonstrates that this shift was not permanent.

D. Implications

Based on the parameter estimates for consumption data in Table I, we solve the model numerically to study its properties.¹⁰ In the following, we are interested in how the perception

⁹The likelihood ratio test statistic is 2.15 which is χ^2 -distributed with 8 degrees of freedom. The corresponding critical value at 10% is 13.36.

¹⁰More details on the solution procedure are contained in the appendix. There we also report model implied moments for stock returns and risk-free rate.

about the conditional moments of consumption growth affect the wealth-consumption ratio. To this end, we define the perceived first and second moments of consumption growth as belief weighted averages

$$\hat{\mu}_t = b_{\mu,t}\mu_h + (1 - b_{\mu,t})\mu_l \quad \hat{\sigma}_t = b_{\sigma,t}\sigma_h + (1 - b_{\sigma,t})\sigma_l \quad (12)$$

and the corresponding changes in the perceived moments as

$$\Delta\hat{\mu}_t = \hat{\mu}_t - \hat{\mu}_{t-1} \quad \Delta\hat{\sigma}_t = \hat{\sigma}_t - \hat{\sigma}_{t-1} \quad (13)$$

In Figures 2 and 3, we plot the wealth-consumption ratio as a function of the perceived conditional first $\hat{\mu}_t$ (left graph) and second $\hat{\sigma}_t$ (right graph) moments of consumption growth when the agent has a high EIS of 1.5 (Figure 2) and a low EIS of 0.5 (Figure 3). We further calibrate the model to a quarterly rate of time preference, β , of 0.995 and relative risk aversion, γ , of 30. Risk aversion of 30 seems unrealistically high. This section, however, is meant to yield qualitative guidance and not quantitative results. Figure 2 illustrates that 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. The opposite is true when the EIS equals 0.5 as in Figure 3.

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., $PC = 1/(r_f + r_E - g)$. 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 one when the EIS is greater than one.

In order to test the model in the cross-section of returns, it is convenient to restate the fundamental asset pricing equation (7) in terms of betas¹¹

$$\mathbb{E}_t[R_{i,t+1}^e] \approx \beta_{c,t}^i \lambda_{c,t} + \beta_{\mu,t}^i \lambda_{\mu,t} + \beta_{\sigma,t}^i \lambda_{\sigma,t} \quad (14)$$

where $\beta_{c,t}^i, \beta_{\mu,t}^i, \beta_{\sigma,t}^i$ denote risk loadings of asset i at date t with respect to consumption growth, and changes in the conditional first and second moments of consumption growth and $\lambda_{c,t}, \lambda_{\mu,t}, \lambda_{\sigma,t}$ are the respective market prices of risk given by

$$\lambda_{c,t} = \gamma \text{Var}_t(\Delta c_{t+1}) \quad \lambda_{\mu,t} = A(1 - \theta) \text{Var}_t(\Delta \hat{\mu}_{t+1}) \quad \lambda_{\sigma,t} = B(1 - \theta) \text{Var}_t(\Delta \hat{\sigma}_{t+1}) \quad (15)$$

where A and B are the sensitivities of the wealth-consumption ratio with respect to changes in the conditional first and second moments of consumption growth. The main cross-sectional implications of the model are the following. Assuming that the EIS is greater than the inverse of relative risk aversion ($\psi > 1/\gamma$), the agent requires lower expected excess returns for stocks which load less (low betas) on expected consumption growth and more (high betas) on consumption growth volatility.

Even though the sign switch of the sensitivity coefficients A and B occurs at unity, as explained above, the market prices of the conditional growth rate and volatility of consumption

¹¹To derive Equation (14), we have to assume that the log wealth-consumption is approximately affine in the perceived first and second moments of consumption growth implying that $\Delta p_{ct} \approx A \Delta \hat{\mu}_t + B \Delta \hat{\sigma}_t$. In the appendix we show that this approximation works well. In particular, we run time-series regressions of the log wealth-consumption ratio on the perceived conditional mean and volatility of consumption growth using simulated data. Even with risk aversion as high as 30, the regression R^2 exceeds 99%.

growth switch sign when the EIS equals the inverse of the coefficient of relative risk aversion ($\psi = 1/\gamma$). When the EIS is greater than the inverse of relative risk aversion ($\psi > 1/\gamma$), the agent prefers intertemporal risk due to the unobservable Markov states to be resolved sooner rather than later. Consequently, she dislikes negative shocks to expected consumption growth and requires a positive market price of risk. At the same time, she likes negative shocks to the conditional volatility of consumption growth and requires a negative market price of risk.

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. Consequently, 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.

These implications do 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 a GARCH consumption model does not give rise to a priced risk factor in a log-linearized approximation to an equilibrium model.¹² Specifically, Equation (4.5) in Restoy and Weil (2004) states that the covariance of any stock with the wealth-consumption ratio is proportional 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 on p. 44: “This is an important result because it embodies the fundamental insight that, for our AR(1)-GARCH(1,1) process, returns are only able to predict future con-

¹²In empirical tests of equilibrium models, GARCH-inspired processes have been used by Adrian and Rosenberg (2008) and Tedongap (2007) to motivate additional factors in the cross-section.

ditional means of consumption growth but carry no information about the future conditional variances.”

To find evidence regarding the magnitude of the representative agent’s EIS, we perform three empirical exercises. First, we estimate time-varying risk loadings on the conditional first and second moments of consumption growth at the firm level and form portfolios based on these loadings. If the agent is not indifferent to intertemporal risk, we expect to find systematic return differences across portfolios. Second, we estimate the market prices directly using portfolios. Both exercises are closely related and we expect findings to be consistent. Third, we run time-series regressions of future excess returns on the perceived first and second moments of consumption growth. The last exercise provides a test whether the EIS is smaller or greater than unity because this relation depends only on the sensitivity of the wealth-consumption ratio with respect to consumption growth moments.

II. Cross-Sectional Return Predictability

The goal of this section is to demonstrate that loadings on the estimated 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 future returns are strongly and robustly predicted by exposure to innovations in consumption volatility, while exposure to consumption growth and changes in expected consumption growth do not help to predict the cross-section of asset prices.

A. Data

Our sample consists of all common stocks ($\text{shrcd} = 10$ or 11) on CRSP that are traded on the NYSE or AMEX ($\text{exchcd} = 1$ 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 2008.

B. Risk Loadings

Our first set of empirical results is based on time-series regressions of individual securities onto log consumption growth and the perceived conditional mean and volatility of consumption growth. In particular, for each security, we estimate factor loadings in each quarter t^* using the previous 20 quarterly observations from

$$R_t^i - R_t^f = \alpha_{t^*}^i + \beta_{c,t^*}^i \Delta c_t + \beta_{\mu,t^*}^i \Delta \hat{\mu}_t + \beta_{\sigma,t^*}^i \Delta \hat{\sigma}_t + \epsilon_t^i \quad (16)$$

where R_t^f denotes the risk-free rate and Δc_t consumption growth for $t \in \{t^* - 19, t^*\}$. Further, $\Delta \hat{\mu}_t$ and $\Delta \hat{\sigma}_t$ are changes in the perceived conditional moments of consumption growth as defined in Equation (13).

The estimated parameters from Equation (16) allow us to evaluate the cross-sectional predictive power of these loadings in two different ways. First, we form portfolios based on the estimated risk exposures and analyze their properties in the time-series. Second, we use cross-sectional regressions as in Fama and MacBeth (1973) to investigate whether the factor loadings help to predict cross-sectional variation in returns.

C. Portfolio Sorts

We now investigate the predictive power of the estimated loadings from model (16) by forming portfolios. This approach has an important advantage relative to Fama-McBeth regressions where a potential error-in-variable problem leads to underestimated standard errors. In contrast, statistical inference based on portfolios is conservative. When variables are measured

with noise, the portfolio assignment will be less accurate as some stocks are sorted into the wrong group. 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.

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 (16). Table II 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 consumption growth, $\beta_{c,t}^i$. Consistent with prior research (e.g., Mankiw and Shapiro (1986) and Lettau and Ludvigson (2001b)), an asset's contemporaneous short horizon loading on consumption growth does not help to generate a return differential across portfolio for either weighting scheme. A similar result follows by forming portfolios based on changes in beliefs about expected consumption growth, $\beta_{\mu,t}^i$ (Panel B). In contrast, exposure to consumption volatility risk, $\beta_{\sigma,t}^i$, predicts future returns strongly and negatively (Panel C). Stocks that comove highly with changes in consumption volatility underperform their peers in the future. An equally-weighted strategy results in a return of the long-short portfolio of -0.19% monthly. The value weighted return is even larger (in absolute value) with -0.43% per month or in excess of -5% annually. In Panel D, we repeat the analysis but we control for market returns in the time-series estimation of risk loadings. By comparing Panels C and D, we observe that all point estimates are nearly identical but the t -statistics on the zero cost portfolio are now larger. By including the market return, consumption volatility risk loadings have a purely cross-sectional interpretation since the market controls for time-series variation not captured by consumption growth. Consequently, standard errors are smaller. We focus on this specification in the remaining cross-sectional analysis.

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 the EIS is greater than the inverse of RRA, the model predicts that returns are positively related to $\beta_{\mu,t}$ and negatively to $\beta_{\sigma,t}$. We do not find convincing evidence that exposure to fluctuations in expected consumption growth predicts returns but exposure to fluctuations in consumption volatility does so negatively. This finding is consistent with the model only if the agent dislikes intertemporal risk and the EIS is greater than the inverse of RRA.

D. Robustness

Cross-sectional differences in returns might not be surprising if consumption volatility betas covary with other variables known to predict returns. Crucially, Table III 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 quintiles 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 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 Table IV, stocks are independently sorted into three portfolios based on $\beta_{\sigma,t}^i$, and into two portfolios based on market capitalization (Panel A) or book-to-market ratio

(Panel B). 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 Panel A shows that consumption volatility risk is consistently present and strong for both equal and value-weighted strategies with return differences ranging from -0.09% to -0.23% monthly. 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, Panel B confirms that consumption volatility risk is also present within book-to-market groups.

As an alternative to portfolio sorts, we also perform Fama-MacBeth regressions by cross-sectionally regressing monthly returns of each asset onto its latest available risk loadings as well as size and value characteristics. The explanatory variables are normalized each quarter so they are centered around zero with unit variance. 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.

The results of the Fama-MacBeth regressions are presented in Table V. Model specifications I-III present univariate effects of each risk loading. Confirming previous findings, the average coefficients on consumption growth betas, $\beta_{c,t}^i$, as well as expected consumption growth betas, $\beta_{\mu,t}^i$, are small and insignificant. Exposure to consumption volatility risk, however, as measured by $\beta_{\sigma,t}^i$, shows up strongly negative and significant. Specification IV is the full model. Now, both the loading on consumption growth and consumption volatility risk are significant. In regression V, we add two characteristics known to predict stock returns, namely, the market capitalization (ME_t^i) and the ratio of book value of equity to market value (BM_t^i), to confirm that the predictive power of consumption volatility is not already

captured by these predictors. The absolute value of the point estimate is slightly reduced by the addition of the two characteristics, but it remains significant.

III. 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. We find 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. Alternatively, we also form a long-short portfolio based on consumption volatility risk (CVR) and demonstrate that it shows up strongly and significantly as a priced factor in cross-sectional regressions. While the CVR portfolio only modestly correlates with the value factor 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 (14) states that, in a log-linear approximation, expected excess returns depend on consumption growth and changes of the perceived conditional first and second moments of consumption growth. We evaluate the performance of our model in two stages. First, for each test asset, we obtain risk loadings from the time-series regression

$$R_t^i - R_t^f = \alpha^i + \beta_c^i \Delta c_t + \beta_\mu^i \Delta \hat{\mu}_t + \beta_\sigma^i \Delta \hat{\sigma}_t + \epsilon_t^i \quad (17)$$

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 VI. 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 , mean absolute pricing error (MAPE) and 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 ($\Delta \hat{\mu}_t$), as well as changes in beliefs about consumption growth volatility ($\Delta \hat{\sigma}_t$). 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 CVR.

As test assets, we use the 25 Fama-French portfolios in Panel A which have been shown to challenge the single factor CAPM. Lewellen, Nagel, and Shanken (2008) criticize the use of only those 25 portfolios as test assets since they exhibit a strong factor structure. Following their suggestions, we also expand the set of assets. In Panel B, we add the 5 value-weighted consumption volatility risk portfolios (Table II, Panel D) as test assets. To broaden the scope beyond equity pricing, we also consider the 6 CRSP bond return portfolios with maturities of 1, 2, 3, 4, 5, and 10 years in addition to the 5 volatility risk portfolios in Panel C.

Regression I in each panel shows results for the standard consumption CAPM. Confirming prior research, the market price of consumption risk in Panels A and B is insignificant and low R^2 s indicate that the C-CAPM performs poorly in pricing the set of test assets. When bond returns are included as test assets (Panel C), the C-CAPM performs better because there is a large spread in returns and betas between asset classes. Regression II in each panel reports the full three factor model (17). Similar to our previous findings, the market price of expected consumption growth, $\Delta \hat{\mu}_t$, is insignificant. In contrast, consumption volatility risk, $\Delta \hat{\sigma}_t$, is a priced factor in the cross-section independent of the test assets. Importantly, the price of volatility risk is negative which is consistent with our portfolio sort results.

Alternatively, we form a consumption volatility risk (CVR) portfolio as a proxy for $\Delta \hat{\sigma}_t$ to reduce measurement error in consumption volatility.¹³ The CVR 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 $\beta_{\sigma,t}^i$ in Table II, Panel D. We do not form a factor based on loadings

¹³While 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)). Moreover, both $\Delta \hat{\mu}_t$ and $\Delta \hat{\sigma}_t$ are estimates and themselves depend on the imposed model for consumption growth dynamics.

on expected consumption growth, $\Delta\hat{\mu}_t$, 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.

Regression III in each panel of Table VI shows a significantly negative price of risk for the CVR factor, while beliefs about the mean consumption growth continue to be insignificant. Replacing the estimated consumption volatility with a traded portfolio results in drastic improvements in second stage R^2 .

The predictions of our theory in Section I depend on the preference parameters of the representative agent. 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 $\Delta\hat{\sigma}_t$ and its mimicking CVR portfolio are significantly negative, thus suggesting an EIS greater than the inverse of RRA for the representative agent.

Figure 4 displays average second stage pricing errors of quarterly excess returns of the 25 Fama-French portfolios (black dots) and 5 volatility risk portfolios (red stars), as in Table VI, Panel B. Each graph 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 first graph depicts the consumption CAPM (Regression I). Visually, this graph confirms that the consumption CAPM does not perform well in pricing the 30 test portfolios. While the portfolios vary drastically in their average realized returns, the model predicted returns are all very close together, resulting in a narrow cloud. In the second graph, we present the full model (Regression II) and in the third graph, we substitute $\Delta\hat{\sigma}_t$ by its mimicking factor CVR. Both graphs confirm that in the full specifications 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. Even though CVR is independent of the book-to-market characteristic and comoves only modestly with HML, we find that substituting HML with CVR in the Fama-French three factor model results in similar pricing and leaves pricing errors unaffected.

Summary statistics for the CVR portfolio are given in Table VII. The CVR portfolio has a mean return of -0.44% and a standard deviation of 3.40% 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 size factor SMB (0.08) and close to the Sharpe ratio of value factor HML (0.15). The correlation matrix of the pricing factors (Panel B) shows that the CVR portfolio returns are uncorrelated with the market. The correlations with the SMB and HML factors are moderate at 17% and -26% , respectively, even though the CVR portfolio is neutral with respect to size and book-to-market characteristics (see Panel B of Table III). To put these correlations in perspective, we note that all the pairwise correlations between the Fama-French factors are larger. Parameter estimates from a time-series regression of the CVR factor onto the other factors are reported in Panel C. The CAPM (regression II) does not explain the returns of the CVR portfolio. In regression III, the Fama-French factors attenuate the estimated intercept $\hat{\alpha}$ towards zero, but it remains large and significant. This reduction is solely driven by HML and both the market and SMB have insignificant coefficients. The three factors only explain about 8% of the variation in the CVR factor.

Table VIII reports factor loadings from regressions of excess returns on the market excess return ($R_{M,t}^E$) and the CVR factor

$$R_t^i - R_t^f = \alpha^i + \beta_M^i R_{M,t}^E + \beta_{CVR}^i CVR_t + \epsilon_t^i. \quad (18)$$

Panel A reports estimated coefficients for the five value-weighted book-to-market portfolios, and Panel B for our five volatility risk sorted portfolios. We observe that the loadings of the value portfolios on the volatility risk factor decrease from growth to value portfolios. A low

risk exposure is consistent with high expected returns for value stocks since the price of CVR risk is negative. Loadings on the CVR factor therefore suggest a risk based explanation of the value anomaly.

For the volatility risk portfolios, the loading on the CVR factor increases monotonically from -0.56 to 0.44 . This finding can be interpreted as evidence that volatility exposure is a systematic source of risk because the portfolios move together. Cochrane (2001) points out that comovement would not be expected if the return differentials are explained by characteristics.

To relate the pricing implications of CVR to existing factors, Table IX shows estimated prices of risks for the market excess return (MKT), size (SMB), value (HML), and consumption volatility (CVR) risk factors, 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 12 monthly lags. We also report second stage R^2 , mean absolute pricing error (MAPE) and the model J -test (χ^2 statistic) with its associated p -value (in percent). The test assets considered are the 25 Fama-French portfolios augmented with our five consumption volatility risk portfolios.

Regressions I and III show the results for the benchmark models, the market CAPM (I) and the Fama-French model (III). 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 less than 7%. The three factor model reduces the pricing errors significantly and yields an R^2 of 76%. 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 benchmark factors with CVR. In all specifications, the estimates for the price of a unit CVR risk are significant and negative, ranging from -0.45% to -0.56% monthly. These estimates are remarkably close to the mean return of the CVR factor of -0.44% . In regression III, the factors are the market portfolio and CVR. This specification yields improvements over the one factor market model. Interestingly, although CVR is based on consumption data, a three factor model based on the market,

SMB and CVR (regression IV) generates an identical mean absolute pricing error to the Fama-French model. Augmenting the Fama-French three-factor model with CVR as a fourth factor (regression V) leads to a marginal improvement in the model’s ability to price the cross-section. In summary, replacing HML with CVR does not deteriorate the model’s performance, while including both CVR and HML as factors improve the model fit only slightly. Hence, HML and CVR are substitutes in cross-sectional pricing for our test assets. 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 CVR portfolio is based on a Markov model for low-frequency consumption data. Interestingly, their short-run volatility component has similar pricing implications to CVR, whereas their long-run component performs worse than CVR. 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 CVR factor thus has a much different and macroeconomically more meaningful interpretation.

Figure 5 replicates Figure 4 for market based pricing models. The 25 size-value portfolios are represented by black dots, and the five volatility risk portfolios by red stars. The top left graph depicts the CAPM (regression I in Table IX). The remaining graphs show the CAPM augmented with the volatility risk factor (top right graph, regression II), the Fama-French three factor model (bottom left graph, regression III), and a three factor model that uses CVR instead of HML (bottom right graph, regression IV). Visually, these graphs confirm that simply adding CVR to the market factor improves the model fit. At first sight, both the Fama-French model and the three factor CVR model seem to price the 30 portfolios well. Upon closer inspection, however, the Fama-French model does not succeed in generating a spread in predicted excess returns of the five volatility risk portfolios as indicated by the narrow cloud of red stars. In contrast, the three-factor CVR model works well for both size-value portfolios and consumption volatility risk portfolios.

IV. Time Series Predictability

In the previous sections, we have shown that loadings on consumption growth volatility predict returns cross-sectionally and that 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 I, the model implies a negative relation between expected returns and expected consumption growth and a positive relation between expected returns and consumption growth volatility when the EIS is greater than unity. Noting that the wealth-consumption ratio is inversely related to expected returns, this effect can be seen in Figure 2. The opposite holds when the EIS is smaller than unity (see Figure 3).

Table X reports forecasting regressions of quarterly excess market returns onto lagged variables. The data ranges from 1955 through 2008. Predictors are the dividend yield (DY), payout ratio (DE), the term spread (TS), aggregate book-to-market ratio (BM), the consumption-wealth ratio of Lettau and Ludvigson (2001a) (*cay*), and changes in beliefs about the first ($\Delta\hat{\mu}_t$) and second moments ($\Delta\hat{\sigma}_t$) of consumption growth.¹⁴ All of those have been shown to predict stock returns at various horizons.¹⁵ 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.

Regressions I-III show the benchmark results of multivariate predictive regressions. The four standard predictor variables jointly result in an R^2 of about 5%. 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 beliefs about the volatility state show up economically and statistically significant and yield a regression R^2 of 2.6% in a univariate regression. The R^2 of *cay* in the univariate regression II is somewhat larger at 4%. The economic impact of consumption volatility risk is large. A one standard deviation increase in

¹⁴We thank Amit Goyal and Martin Lettau for making their data available.

¹⁵See, for example, Fama and Schwert (1977), Rozeff (1984), Keim and Stambaugh (1986), Campbell (1987), Campbell and Shiller (1988), Lamont (1998), Campbell and Thompson (2008), and Goyal and Welch (2008).

$\Delta\hat{\sigma}_t$ results in an increase in the expected risk premium of 1.4% quarterly.¹⁶ The economic impact is similar to *cay* (1.6% quarterly).

Regressions VI to VIII demonstrate that the marginal impact of $\Delta\hat{\sigma}_t$ remains strong and significant even after controlling for all other predictors, including *cay*. Moreover, the coefficients on consumption volatility are virtually unaffected by the inclusion of other predictors, indicating that its forecasting ability is orthogonal to existing variables. The predictive R^2 exceeds 9% in the multivariate setting with all predictor variables.

The observation that consumption volatility and *cay* are orthogonal is surprising. In our model, changes in consumption volatility enter the pricing kernel only because they affect the wealth-consumption ratio. Thus, one might expect that direct measures of the wealth-consumption ratio, such as *cay*, comprise all relevant information about the volatility state. Our findings therefore suggest 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, for instance, 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. For the variable $\Delta\hat{\sigma}_t$, this bias is less of a concern since it is not a price scaled variable. The appendix shows this bias is immaterial in our setup.

We acknowledge that the predictive results presented have limitations. First, they are in sample results. Second, there is a look-ahead bias in $\Delta\hat{\sigma}_t$. 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. 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

¹⁶Note that $\Delta\sigma_t$ has a standard deviation of around 0.00035.

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.

V. 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 follow Lettau, Ludvigson, and Wachter (2008), who 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 (1991) 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. In the cross-section, we first show that firm level loadings on changes in beliefs about consumption volatility significantly forecast returns, while loadings on changes in beliefs about expected consumption growth do not. A negative relation between betas and future returns indicates a negative price of consumption volatility risk. This is confirmed in cross-sectional pricing tests, where both consumption volatility and its mimicking portfolios are negatively priced sources of risk. In the context of our model, these findings suggest an EIS greater than the inverse of the RRA for the representative agent.

In time-series tests, we find that shocks to beliefs about the volatility state forecast the equity premium. In a univariate regression, changes about perceived consumption volatility achieve an R^2 of 2.6%. A one standard deviation increase in perceived volatility is followed by an increase of the equity returns of 1.4% quarterly. The economic impact is comparable to the

one generated by *cay* of Lettau and Ludvigson (2001a). Surprisingly, *cay*, a direct measure of the wealth-consumption ratio, does not subsume the predictive power of consumption volatility. The positive coefficient of consumption volatility in the predictive regressions indicates that the representative agent has elasticity of intertemporal substitution greater than unity.

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.
- Ang, Andrew, Jun Liu, and Krista Schwarz, 2008, Using stocks or portfolios in tests of factor models, Columbia University Working Paper.
- 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, Dana Kiku, and Amir Yaron, 2007, Risks for the long run: Estimation and inference, Working Paper.
- 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-fitted 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.
- Cochrane, John H., 2001, *Asset Pricing* (Princeton University Press: Princeton, N.J.).
- 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 vol 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.
- Ferson, Wayne E., and Campbell R. Harvey, 1992, Seasonality and consumption-based asset pricing, *Journal of Finance* 47, 511–552.

- French, Kenneth R., G. William Schwert, and Robert F. Stambaugh, 1987, Expected stock returns and volatility, *Journal of Financial Economics* 19, 3–29.
- 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.
- Goyal, Amit, and Ivo Welch, 2008, A comprehensive look at the empirical performance of equity premium prediction, *Review of Financial Studies* 21, 1455–1508.
- 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.
- Lamont, Owen A., 1998, Earnings and expected returns, *Journal of Finance* 53, 1563–1587.
- 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.
- , Stefan Nagel, and Jay Shanken, 2008, A skeptical appraisal of asset pricing tests, *Journal of Financial Economics*.
- 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, inflation, 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 fluctuations, *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.
- Rozeff, 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.
- 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 processes for being in the high expected growth rate state (top figure) and high volatility state (bottom figure). The estimation procedure follows Hamilton (1994). We use quarterly per capita real consumption expenditure for non-durable goods and services for the years 1955.Q1-2008.Q4.




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) moments of consumption growth. The dynamics of the underlying consumption process are summarized in Table I. The representative agent has an EIS of 1.5, RRA of 30 and quarterly rate of time preference of 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) moments of consumption growth. The dynamics of the underlying consumption process are summarized in Table I. The representative agent has an EIS of 0.5, RRA of 30 and quarterly rate of time preference of 0.995.




Figure 4. Pricing Errors of the Consumption-Based Model

This figure depicts average quarterly excess returns of the 25 Fama-French portfolios (black dots) and 5 volatility risk portfolios (red stars) against model predicted excess returns. The first graph represents the standard consumption CAPM and the second graph the full model with consumption growth (Δc_t) as well as changes in the perceived first ($\Delta \hat{\mu}_t$) and second ($\Delta \hat{\sigma}_t$) moments of consumption growth as explanatory factors. In the third graph, we replace beliefs about the volatility state with the CVR 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 (black dots) and 5 volatility risk portfolios (red stars) 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 CVR factor to the CAPM and, in the bottom-right graph, we replace the HML factor with the CVR 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

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

where $\mu_t \in \{\mu_l, \mu_h\}$ and $\sigma_t \in \{\sigma_l, \sigma_h\}$ follow independent Markov processes with transition matrices P^μ and P^σ , respectively,

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

The estimation procedure follows Hamilton (1994). We use quarterly per capita real consumption expenditure for non-durable goods and services for the years 1955.Q1-2008.Q4. Standard errors are reported in parentheses.

Panel A: Parameter Estimates (%)			
μ_l	μ_h	σ_l	σ_h
0.3668	0.7800	0.2112	0.4772
(0.0376)	(0.0600)	(0.0279)	(0.0532)
Panel B: Marginal Transition Probabilities			
p_μ^{ll}	p_μ^{hh}	p_σ^{ll}	p_σ^{hh}
0.9304	0.8929	0.9452	0.9617
(0.0353)	(0.0597)	(0.0400)	(0.0458)

Table II
Portfolios Formed on Risk Exposure

This table reports average equally-weighted (EW) and value-weighted (VW) monthly returns (%) of portfolios based on time-varying loadings from rolling time-series regressions (16). Specifically, we regress individual returns on log consumption growth ($\beta_{c,t}^i$) and changes in the perceived first ($\beta_{\mu,t}^i$) and second moments ($\beta_{\sigma,t}^i$) of consumption growth using 20 quarterly observations. In Panel D, we also control for fluctuations in the market return when estimating risk loadings. The t -statistics reported in parentheses are based on Newey-West adjusted standard errors using 12 lags. The sample period is from January 1964 to December 2008.

Panel A: Univariate Sorts Based on $\beta_{c,t}^i$						
	Low		Med		High	High - Low
EW	1.09 (4.01)	1.16 (5.19)	1.15 (5.09)	1.17 (4.68)	1.25 (4.16)	0.16 (1.16)
VW	0.84 (3.62)	0.89 (4.50)	0.84 (4.25)	0.88 (4.53)	0.91 (3.69)	0.07 (0.37)

Panel B: Univariate Sorts Based on $\beta_{\mu,t}^i$						
	Low		Med		High	High - Low
EW	1.22 (4.42)	1.17 (5.13)	1.14 (5.03)	1.14 (4.82)	1.15 (3.86)	-0.06 (-0.57)
VW	0.95 (4.32)	0.90 (4.72)	0.81 (4.08)	0.86 (4.17)	0.91 (3.38)	-0.03 (-0.22)

Panel C: Univariate Sorts Based on $\beta_{\sigma,t}^i$						
	Low		Med		High	High - Low
EW	1.28 (4.26)	1.16 (4.77)	1.15 (5.02)	1.16 (5.14)	1.07 (4.00)	-0.19 (-1.68)
VW	1.08 (4.26)	0.90 (4.42)	0.85 (4.30)	0.89 (4.62)	0.65 (2.73)	-0.43 (-2.61)

Panel D: Univariate Sorts Based on $\beta_{\sigma,t}^i$ - Controlling for R_M						
	Low		Med		High	High - Low
EW	1.30 (4.46)	1.19 (5.16)	1.17 (5.17)	1.08 (4.71)	1.08 (3.79)	-0.21 (-2.09)
VW	1.08 (4.42)	0.93 (4.93)	0.82 (4.21)	0.89 (4.52)	0.64 (2.71)	-0.44 (-3.10)

Table III
Characteristics of Consumption Volatility Risk Portfolios

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

Panel A: Univariate Sorts Based on $\beta_{\sigma,t}^i$						
	Low	Med	High	High - Low		
Return	1.08	0.93	0.82	0.89	0.64	-0.44
$\beta_\sigma/100$	-2.69	-0.87	-0.03	0.84	2.93	5.62
Panel B: Characteristics of Sorts Based on $\beta_{\sigma,t}^i$						
	Low	Med	High	High - Low		
ME	1428.21	2418.57	2388.74	1742.64	892.62	-535.59
ME Rank	4.25	5.43	5.64	5.28	4.13	-0.12
BM	0.91	0.92	0.91	0.91	0.90	-0.01
BM Rank	4.87	5.02	5.04	4.99	4.83	-0.04

Table IV
Portfolios Formed on Consumption Volatility Risk and Characteristics

This table reports average equally-weighted and value-weighted monthly returns (%) of independent double sorts based on consumption volatility loadings ($\beta_{\sigma,t}^i$) and market capitalizations in Panel A and based on consumption volatility loadings and book-to-market ratios in Panel B. The t -statistics reported in parentheses are based on Newey-West adjusted standard errors using 12 lags.

Panel A: Portfolios Formed on Consumption Volatility Risk and Market Capitalization								
	Equally-Weighted Returns				Value-Weighted Returns			
	Low	Med	High	H - L	Low	Med	High	H - L
Small	1.36	1.36	1.26	-0.09	1.26	1.28	1.10	-0.15
	(4.43)	(4.84)	(4.13)	(-1.25)	(4.60)	(5.06)	(3.93)	(-1.79)
Big	1.05	1.04	0.90	-0.15	0.96	0.85	0.73	-0.23
	(4.61)	(5.15)	(3.96)	(-1.72)	(4.68)	(4.54)	(3.40)	(-2.03)
S - B	-0.30	-0.31	-0.35		-0.29	-0.42	-0.36	
	(-2.02)	(-2.27)	(-2.34)		(-1.70)	(-2.56)	(-1.97)	
Panel B: Portfolios Formed on Consumption Volatility Risk and Book-to-Market Ratio								
	Equally-Weighted Returns				Value-Weighted Returns			
	Low	Med	High	H - L	Low	Med	High	H - L
Low BM	1.06	1.01	0.89	-0.17	0.90	0.84	0.65	-0.25
	(3.90)	(4.45)	(3.39)	(-1.84)	(4.16)	(4.31)	(2.89)	(-2.08)
High BM	1.44	1.31	1.30	-0.14	1.17	1.00	1.10	-0.06
	(5.42)	(5.60)	(4.77)	(-1.79)	(5.46)	(5.32)	(5.12)	(-0.53)
H - L	0.38	0.30	0.41		0.26	0.16	0.46	
	(3.99)	(3.50)	(3.89)		(2.10)	(1.37)	(3.30)	

Table V
Fama-MacBeth Regressions

This table reports cross-sectional regressions of monthly returns on lagged 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 of log consumption growth ($\beta_{c,t}^i$), changes in the perceived conditional mean ($\beta_{\mu,t}^i$) and volatility ($\beta_{\sigma,t}^i$) 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 to the test month. All explanatory variables are normalized so they are centered around zero with unit variance. We report time-series averages of the second stage coefficients. The t -statistics in parentheses are based on Newey-West adjusted standard errors using 12 lags.

	$\beta_{c,t}^i$	$\beta_{\mu,t}^i$	$\beta_{\sigma,t}^i$	ME_t^i	BM_t^i
I	0.05 (0.97)				
II		-0.02 (-0.54)			
III			-0.08 (-2.22)		
IV	0.16 (2.03)	0.02 (0.39)	-0.09 (-1.62)		
V			-0.06 (-1.79)	-0.07 (-1.43)	0.18 (3.98)

Table VI
Volatility Risk Pricing

This table reports market prices of risk from two stage regressions where Δc_t denotes log consumption growth, $\Delta \hat{\mu}_t$ and $\Delta \hat{\sigma}_t$ are changes in filtered beliefs about the first and second moments of consumption growth. The CVR factor is the return of holding a long position in the value-weighted quintile of stocks with high volatility risk ($\beta_{\sigma,t}^i$) and a short position in low volatility risk, as reported in Panel D of Table II. In Panel A, the test assets are the value-weighted 25 Fama-French value and size portfolios, in Panel B, we add the 5 value-weighted consumption volatility risk portfolios as in Panel D of Table II, and in Panel C, we add 6 bond portfolios from CRSP. The data covers January 1964 to December 2008. The t -statistics are corrected for estimation error in the first stage as proposed by Shanken (1992) and are Newey-West adjusted using 4 lags. For each specification, we report the R^2 , mean absolute pricing error (MAPE) in parentheses, regression J -statistic (χ^2) with the associated p -value (in %).

	Const. (t -stat)	Δc (t -stat)	$\Delta \hat{\mu}$ (t -stat)	$\Delta \hat{\sigma}$ (t -stat)	CVR (t -stat)	R^2 (MAPE)	χ^2 (p -val)
Panel A: 25 Value-Size Portfolios							
I	1.54 (2.01)	0.10 (0.52)				2.54 (0.62)	92.44 (0.00)
II	1.33 (1.31)	0.25 (1.29)	-0.02 (-0.57)	-0.03 (-2.26)		9.91 (0.60)	46.18 (0.19)
III	0.79 (0.64)	0.01 (0.03)	0.03 (0.94)		-6.49 (-2.46)	75.47 (0.31)	28.13 (17.13)
Panel B: 25 Value-Size Portfolios & 5 CVR Portfolios							
I	1.13 (1.46)	0.16 (0.83)				6.11 (0.60)	122.33 (0.00)
II	0.94 (0.79)	0.30 (1.35)	-0.02 (-0.49)	-0.04 (-2.24)		13.83 (0.58)	46.62 (1.09)
III	1.06 (1.25)	0.06 (0.42)	0.02 (0.80)		-2.84 (-3.99)	42.35 (0.46)	75.00 (0.00)
Panel C: 25 Value-Size Portfolios & 5 CVR Portfolios & 6 Bond Portfolios							
I	0.69 (2.84)	0.25 (1.87)				42.56 (0.52)	146.40 (0.00)
II	0.87 (1.77)	0.29 (1.86)	-0.01 (-0.38)	-0.04 (-2.30)		47.05 (0.52)	67.11 (0.04)
III	0.43 (1.18)	0.24 (1.91)	0.02 (1.01)		-2.90 (-4.07)	63.12 (0.41)	98.87 (0.00)

Table VII
Volatility Risk Factor

This table provides descriptive statistics of the volatility risk (CVR) portfolio. The CVR portfolio is the return of holding a long position in the value-weighted quintile of stocks with high volatility risk ($\beta_{\sigma,t}^i$) and a short position in low volatility risk, as reported in Panel D of Table II. Panel A shows the mean, standard deviation and the Sharpe Ratio of the CVR portfolio as well as the three Fama-French factors. Panel B presents the correlation matrix of the factor returns. Panel C reports parameter estimates from time-series regressions of the CVR portfolio on the market and the three Fama-French factors with Newey-West adjusted standard errors using 12 lags.

Panel A: Summary Statistics					
	CVR	MKT	SMB	HML	
Mean (%)	-0.44	0.37	0.26	0.43	
Std. Dev. (%)	3.40	4.47	3.20	2.90	
Sharpe Ratio	-0.13	0.08	0.08	0.15	
Panel B: Correlations					
	CVR	MKT	SMB	HML	
MKT	0.09	1			
SMB	0.17	0.30	1		
HML	-0.26	-0.38	-0.26	1	
Panel C: Time-Series Regressions					
	α (%)	β_{MKT}	β_{SMB}	β_{HML}	R^2 (%)
I	-0.44 (-3.10)				
II	-0.47 (-3.21)	0.07 (1.27)			0.86
III	-0.34 (-2.35)	-0.03 (-0.49)	0.12 (0.87)	-0.29 (-2.51)	7.92

Table VIII
Factor Exposures to the CVR Factor

This table reports factor loadings of the 5 value portfolios (Panel A) and 5 volatility risk portfolios (Panel B) with the market return (MKT) and consumption volatility risk factor (CVR). The CVR portfolio is the return of holding a long position in the value-weighted quintile of stocks with high volatility risk ($\hat{\beta}_{\sigma,t}^i$) and a short position in low volatility risk, as reported in Panel D of Table II. The data starts in January 1964 and ends in December 2008. The t -statistics reported in parentheses are based on Newey-West adjusted standard errors using 12 lags.

	Low	Med	High		
Panel A: Value Portfolios					
α (%)	-0.09 (-1.28)	0.02 (0.30)	0.10 (1.17)	0.22 (2.22)	0.33 (2.86)
β_{MKT}	1.06 (67.89)	0.99 (38.02)	0.91 (31.86)	0.86 (24.64)	0.95 (20.85)
β_{CVR}	-0.01 (-0.41)	-0.09 (-2.02)	-0.06 (-1.31)	-0.14 (-2.66)	-0.16 (-2.67)
Panel B: Volatility Risk Portfolios					
α (%)	-0.02 (-0.38)	0.03 (0.36)	-0.02 (-0.28)	0.09 (1.27)	-0.02 (-0.38)
β_{MKT}	1.08 (40.13)	0.89 (30.30)	0.88 (37.25)	0.96 (48.71)	1.08 (40.13)
β_{CVR}	-0.56 (-16.84)	-0.26 (-4.94)	-0.11 (-2.21)	0.04 (0.52)	0.44 (13.21)

Table IX
Volatility Risk Pricing Factor

This table reports market prices of risk from two stage regressions where MKT denotes the market excess return, and SMB and HML are the Fama-French size and value factors. The CVR factor is the return of holding a long position in the value-weighted quintile of stocks with high volatility risk ($\beta_{\sigma,t}^i$) and a short position in low volatility risk, as reported in Panel D of Table II. Test assets are the value-weighted 25 Fama-French size and book-to-market portfolios as well as the 5 value-weighted consumption volatility risk portfolios as in Panel D of Table II. The t -statistics are corrected for estimation error in the first stage as proposed by Shanken (1992) and are Newey-West adjusted with 12 lags to account for heteroskedasticity and autocorrelation. For each specification, we report the R^2 , mean absolute pricing error (MAPE) in parentheses, regression J -statistic (χ^2) with the associated p -value (in %). The data covers January 1964 to December 2008.

	Const.	MKT	SMB	HML	CVR	R^2	χ^2
	(t -stat)	(t -stat)	(t -stat)	(t -stat)	(t -stat)	(MAPE)	(p -val)
I	1.00 (2.47)	-0.39 (-0.85)				6.73 (0.20)	112.31 (0.00)
II	0.60 (1.60)	-0.03 (-0.07)			-0.52 (-3.30)	15.09 (0.20)	99.02 (0.00)
III	1.01 (3.21)	-0.62 (-1.68)	0.21 (1.40)	0.48 (3.13)		76.05 (0.09)	82.75 (0.00)
IV	1.75 (3.94)	-1.32 (-2.84)	0.20 (1.30)		-0.56 (-3.51)	74.26 (0.09)	68.59 (0.00)
V	1.00 (3.16)	-0.62 (-1.67)	0.22 (1.42)	0.46 (2.98)	-0.45 (-3.07)	80.06 (0.08)	69.69 (0.00)

Table X
Market Predictability in the Time-Series

This table reports time-series regressions of the market excess return on lagged predictor variables. The market return is the value-weighted CRSP index less the 90 day T-Bill rate. Predictor variables are the dividend yield (DY), payout ratio (DE), the term spread (TS), aggregate book-to-market ratio (BM), the consumption-wealth ratio of Lettau and Ludvigson (2001a) (*cay*), and changes in beliefs about the first ($\Delta\hat{\mu}_t$) and second moments ($\Delta\hat{\sigma}_t$) of consumption growth. The sample period includes the first quarter of 1955 until the fourth quarter of 2008. *t*-statistics are reported in parentheses and based on Newey-West adjusted standard errors using 4 lags.

	Const. (<i>t</i> -stat)	DY (<i>t</i> -stat)	DE (<i>t</i> -stat)	TS (<i>t</i> -stat)	BM (<i>t</i> -stat)	<i>cay</i> (<i>t</i> -stat)	$\Delta\hat{\mu}$ (<i>t</i> -stat)	$\Delta\hat{\sigma}$ (<i>t</i> -stat)	$R^2(\%)$
I	-0.01 (-0.43)	3.83 (2.84)	-0.09 (-1.56)	0.80 (2.08)	-0.12 (-2.06)				5.10
II	0.01 (2.42)					1.13 (3.23)			4.05
III	0.02 (0.57)	1.51 (0.83)	-0.11 (-1.88)	0.60 (1.39)	-0.02 (-0.27)	1.05 (2.01)			6.88
IV	0.01 (2.29)						2.79 (0.29)	37.48 (2.32)	2.64
V	0.01 (2.27)							37.88 (2.40)	2.60
VI	-0.02 (-0.52)	3.67 (2.86)	-0.08 (-1.42)	0.86 (2.24)	-0.12 (-2.00)			38.33 (2.53)	7.74
VII	0.01 (2.34)					1.15 (3.29)		38.84 (2.58)	6.78
VIII	0.02 (0.50)	1.27 (0.74)	-0.10 (-1.75)	0.66 (1.50)	-0.01 (-0.14)	1.08 (2.07)		39.11 (2.69)	9.63

Appendix to Consumption Volatility Risk

A. Wealth-Consumption Ratio Approximation

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

$$PC_t^\theta = \mathbb{E}_t \left[\beta^\theta \left(\frac{C_{t+1}}{C_t} \right)^{1-\gamma} (PC_{t+1} + 1)^\theta \right] \quad (19)$$

Based on the law of iterated expectations, equation (19) can be written as

$$PC_t^\theta = \sum_{i=1}^4 \xi_{t+1|t}(i) PC_{t,i}^\theta \quad (20)$$

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

$$PC_{t,i}^\theta = \mathbb{E} \left[\beta^\theta \left(\frac{C_{t+1}}{C_t} \right)^{1-\gamma} (PC_{t+1} + 1)^\theta \middle| s_{t+1} = i, \xi_{t+1|t} \right] \quad (21)$$

Equation (20) says that the agent forms a belief-weighted average of the state- and belief-conditioned wealth-consumption ratios (21).

Given Equation (20), the local univariate approximations (10) and (11) of the wealth-consumption ratio are derived as follows:

$$\begin{aligned} \Delta pc_{t+1} &= \frac{1}{\theta} \ln \left(\frac{b_{t+1} PC_{t+1,1}^\theta + (1 - b_{t+1}) PC_{t+1,2}^\theta}{b_t PC_{t,1}^\theta + (1 - b_t) PC_{t,2}^\theta} \right) \\ &= \frac{1}{\theta} \ln \left(\frac{(b_t + \Delta b_{t+1}) PC_{t+1,1}^\theta + (1 - (b_t + \Delta b_{t+1})) PC_{t+1,2}^\theta}{b_t PC_{t,1}^\theta + (1 - b_t) PC_{t,2}^\theta} \right) \\ &= \frac{1}{\theta} \ln \left(\frac{b_t PC_{t+1,1}^\theta + (1 - b_t) PC_{t+1,2}^\theta + \Delta b_{t+1} (PC_{t+1,1}^\theta - PC_{t+1,2}^\theta)}{b_t PC_{t,1}^\theta + (1 - b_t) PC_{t,2}^\theta} \right) \\ &= \frac{1}{\theta} \ln \left(1 + \frac{\Delta b_{t+1} (PC_{t+1,1}^\theta - PC_{t+1,2}^\theta)}{b_t PC_{t,1}^\theta + (1 - b_t) PC_{t,2}^\theta} \right) \\ &\approx \frac{1}{\theta} \Delta b_{t+1} \frac{PC_1^\theta - PC_2^\theta}{b_t PC_1^\theta + (1 - b_t) PC_2^\theta} \end{aligned}$$

B. Seasonal Adjustment

The consumption data used in this paper are quarterly, per capita, real consumption of nondurable goods and services, seasonally adjusted at annual rates. Ferson and Harvey (1992)

investigate the asset pricing implications of consumption growth rates obtained from data that are seasonally adjusted with the X-12-ARIMA filter.¹⁷ Ex-ante, the impact of the filter on latent volatility regimes is not obvious.

To measure the impact of the X-12-ARIMA filter, we simulate 300 time series of 200 quarterly log consumption growth rates generated by the Markov model estimated in Table I. We then perturb every fourth quarter data point by +5% and every first quarter data point by -5%, which approximately generates the seasonality observed in the discontinued series for seasonally unadjusted quarterly consumption data. We transform the seasonally perturbed series into consumption levels and apply the X-12-ARIMA filter to get seasonally adjusted consumption data. We lastly estimate the four-state Markov model of Section I on both the original and the seasonally adjusted data.

Table XI shows summary statistics of the estimated Markov chain parameters. We observe that the seasonal adjustment has negligible influence on the estimated states and state-transitions. Moreover, the median correlation between beliefs over states estimated from the original and the filtered data is very high (0.98 for the mean, 0.90 for the standard deviation state). We conclude that the Markov model is robust to the X-12-ARIMA filter.

C. Numerical Solution

Using Equation (20), the wealth-consumption ratio, $PC_t = PC(\xi_{t+1|t})$, solves the following functional equation

$$PC(\xi_{t+1|t}) = \left(\sum_{i=1}^4 \xi_{t+1|t}(i) \mathbb{E} \left[\beta^\theta (PC(\xi_{t+2|t+1}) + 1)^\theta (e^{\mu_i + \sigma_i \epsilon_{t+1}})^{1-\gamma} \mid s_{t+1} = i \right] \right)^{1/\theta}$$

where $\xi_{t+1|t}(i)$ is i -the element of $\xi_{t+1|t}$. We solve this equation as a fixed-point in the wealth-consumption ratio. The grid for the belief state-vector has increments of size 0.025 and the expectation is approximated using Gaus-Hermite quadrature with 10 nodes. Three-dimensional linear interpolation is used between grid points.

¹⁷X-12-ARIMA is a seasonal adjustment program developed at the U.S. Census Bureau. The program is based on the Bureau's earlier X-11 program and the X-11-ARIMA program developed at Statistics Canada.

D. Cross-Sectional Asset Pricing Implications

For our empirical exercise, we assume that the log wealth-consumption ratio is approximately affine in the perceived first and second moments of consumption growth

$$pc_t \approx k + A\hat{\mu}_t + B\hat{\sigma}_t$$

This step provides a more meaningful economic interpretation for mean and volatility states. In Table XII, we confirm the quality of this approximation based on simulations of the model. We simulate 300 economies for 100 years at quarterly frequency. In all three panels, the representative agent has an EIS of 1.5 and a rate of time preference of 0.995. The coefficient of relative risk aversion (RRA) increases from 10 (Panel A) to 20 (Panel B) and 30 (Panel C). In the first regression of each panel, we regress the log wealth-consumption ratio, pc_t , on the prior probabilities of being in a given state, $\xi_{t+1,t}(i)$, $i = 1, 2, 3$. In the second regression, we regress the log wealth-consumption ratio, pc_t , on the perceived first, $\hat{\mu}_t$, and second moment, $\hat{\sigma}_t$, of consumption growth. We report the (across simulation) average regression coefficient and regression R^2 .

Equation (20) 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. The first regression of each panel confirms that the log wealth-consumption ratio is approximately affine in the prior probabilities about the state with regression R^2 exceeding 99%.

The second regression of each panel confirms that the log wealth-consumption ratio is approximately affine in the perceived first and second moments of consumption growth. This approximation captures most variation of changes in the wealth-consumption ratio with regression R^2 exceeding 99%. Intuitively, the third prior probability captures the perceived comovement between the Markov chains for mean and volatility. However, since these two Markov chains are independent by assumption, the third prior probability is redundant.

In order to test the model in the cross-section of returns, it is convenient to restate the

fundamental asset pricing equation (5) in terms of betas

$$\begin{aligned}
\mathbb{E}_t[R_{i,t+1}^e] &\approx -\text{Cov}_t(R_{i,t+1}, m_{t+1}) \\
&= \gamma \text{Cov}_t(R_{i,t+1}, \Delta c_{t+1}) + (1 - \theta) \text{Cov}_t(R_{i,t+1}, \Delta p c_{t+1}) \\
&= \gamma \text{Cov}_t(R_{i,t+1}, \Delta c_{t+1}) + (1 - \theta) A \text{Cov}_t(R_{i,t+1}, \Delta \hat{\mu}_{t+1}) + (1 - \theta) B \text{Cov}_t(R_{i,t+1}, \Delta \hat{\sigma}_{t+1}) \\
&= \beta_{c,t}^i \lambda_{c,t} + \beta_{\mu,t}^i \lambda_{\mu,t} + \beta_{\sigma,t}^i \lambda_{\sigma,t}
\end{aligned}$$

with

$$\beta_{c,t} = \frac{\text{Cov}_t(R_{i,t+1}, \Delta c_{t+1})}{\text{Var}_t(\Delta c_{t+1})} \quad \beta_{\mu,t} = \frac{\text{Cov}_t(R_{i,t+1}, \Delta \hat{\mu}_{t+1})}{\text{Var}_t(\Delta \hat{\mu}_{t+1})} \quad \beta_{\sigma,t} = \frac{\text{Cov}_t(R_{i,t+1}, \Delta \hat{\sigma}_{t+1})}{\text{Var}_t(\Delta \hat{\sigma}_{t+1})}$$

and

$$\lambda_{c,t} = \gamma \text{Var}_t(\Delta c_{t+1}) \quad \lambda_{\mu,t} = A(1 - \theta) \text{Var}_t(\Delta \hat{\mu}_{t+1}) \quad \lambda_{\sigma,t} = B(1 - \theta) \text{Var}_t(\Delta \hat{\sigma}_{t+1})$$

where $\beta_{c,t}^i, \beta_{\mu,t}^i, \beta_{\sigma,t}^i$ denote risk loadings of asset i at date t with respect to consumption growth and the conditional first and second moments of consumption growth, $\lambda_{c,t}, \lambda_{\mu,t}, \lambda_{\sigma,t}$ are the respective market prices of risk.

E. Equity Premium

To quantify the equity premium generated by our model, we first have to specify a process for dividend growth. A common approach is to postulate a levered consumption process for dividends such as $D = C^\lambda$. The Markov switching model allows a more general approach by fitting a Markov model for the conditional first and second moments of dividend growth. Specifically, we assume that log dividend growth follows

$$\Delta d_{t+1} = \mu_t^d + \sigma_t^d \epsilon_{t+1} \quad \epsilon_{t+1} \sim \mathcal{N}(0, 1)$$

where $\mu_t^d \in \{\mu_l^d, \mu_h^d\}$ and $\sigma_t^d \in \{\sigma_l^d, \sigma_h^d\}$ follow the same Markov process as consumption. Consequently, we do not re-estimate the transition matrix of the Markov process but use the estimates reported in Table I. We compute quarterly dividends for the period 1955-2008 using the value-weighted CRSP index with and without distributions. Parameter estimates are summarized in Table XIII.

In Table XIV, we report statistics about the risky and risk-free asset. We simulate 300 economies for 100 years at quarterly frequency. In all three panels, the representative agent has an EIS of 1.5 and rate of time preference of 0.995. The coefficient of relative risk aversion increases from 10 (Panel A) to 20 (Panel B) and 30 (Panel C). We report the average excess return, $\mathbb{E}[R^e]$, the standard deviation of stock returns, $\sigma[R]$, the average risk-free rate, $\mathbb{E}[R^f]$, and the standard deviation of the risk-free rate, $\sigma[R^f]$. In the last row of each panel, we also report moments of the the Markov switching model without learning where the agent knows the state of the economy.

In the specification with RRA of 10, the model generates an annual risk premium of 1%, stock return volatility of 6%, an average risk-free rate of 3% and risk-free rate volatility of 0.3%. This poor performance is not surprising since the Markov chain is not very persistent compared to the specification of Bansal and Yaron (2004). For RRA of 30, the model generates a risk premium of 3.6%.

F. Learning Premium

The Euler equation for the return on wealth is given by

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

Without learning, it follows that

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

since the Markov switches are independent from the normal shocks to consumption growth. In a model with learning, however, the wealth-consumption ratio is correlated with consumption growth, i.e.,

$$\text{Cov}_t \left(\frac{PC_{t+1} + 1}{PC_t}, \frac{C_{t+1}}{C_t} \right) \neq 0$$

since the agent cannot differentiate whether the change in consumption growth comes from the Markov chain or the normal innovation. We call this covariance the learning premium.

To quantify the learning premium, it is convenient to log linearize returns and the pricing kernel. In a log-linear world, the expected excess return on wealth has to satisfy

$$\mathbb{E}_t[R_{w,t+1}^e] \approx -\text{Cov}_t(r_{w,t+1}, m_{t+1})$$

where the log return on wealth can be approximated by

$$r_{w,t+1} \approx \Delta p c_{t+1} + \Delta c_{t+1}$$

and the log pricing kernel by

$$m_{t+1} \approx (\theta \ln \beta - (1 - \theta)(k_0 - k_1 z_t)) - \gamma \Delta c_{t+1} - (1 - \theta) \Delta p c_{t+1}$$

By substituting these two approximation into the expression for the expected excess return on wealth, one obtains

$$\begin{aligned} \mathbb{E}_t[R_{w,t+1}^e] &\approx -\text{Cov}_t(r_{w,t+1}, m_{t+1}) \\ &\approx -\text{Cov}_t(\Delta p c_{t+1} + \Delta c_{t+1}, -\gamma \Delta c_{t+1} - (1 - \theta) \Delta p c_{t+1}) \\ &= \underbrace{\gamma \text{Var}_t(\Delta c_{t+1})}_{\text{short-run}} + \underbrace{(1 - \theta) \text{Var}_t(\Delta p c_{t+1})}_{\text{long-run}} + \underbrace{(1 - \theta + \gamma) \text{Cov}_t(\Delta c_{t+1}, \Delta p c_{t+1})}_{\text{learning}} \end{aligned}$$

Hence, the risk premium in the full model has three components: a short-run, a long-run and a learning premium. The short-run component arises in a model with i.i.d. consumption and power utility. For the long-run component to be non-zero, the model has to contain persistent shocks and the agent has to care about the temporal resolution of risk, i.e., $\theta \neq 1$. The learning premium arises because shocks to consumption growth also lead the agent to update her beliefs about states.

Table XIV can be used to quantify the importance of the learning premium. In the last row of each panel, we also report moments of the the Markov switching model without learning where the agent knows the state of the economy. The difference between the mean excess return generated by the full model and the model without learning is the learning premium. Holding the EIS fixed at 1.5, for RRA of 10 (Panel A), the learning premium is only 7 basis points; for RRA of 20 (Panel B), the learning premium increases to 44 basis points; and for RRA of 30 (Panel C), the learning premium reaches 88 basis points. So the fraction of the total excess return coming from learning increases from 7% to 19% to 24%.

G. Predictive Regression Bias

It is well known that parameter estimates and t -statistics are potentially biased in predictive regressions. Hodrick (1992) shows that using overlapping observations leads to biased inference. More importantly, when the predictor variable is persistent and its innovations are correlated with future returns, Stambaugh (1999), Lewellen (2004), Boudoukh, Richardson, and Whitelaw (2006) and Ang and Bekaert (2007) show that standard econometric techniques can be misleading. 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

$$\begin{aligned}r_t &= \alpha + \beta x_{t-1} + \epsilon_t^r \\x_t &= \phi + \rho x_{t-1} + \epsilon_t^x\end{aligned}$$

where r_t denotes returns and x_t a predictor variable. Lewellen (2004) shows that β estimates are biased by $\gamma(\hat{\rho} - \rho)$ where $\gamma = \text{Cov}(\epsilon_t^r, \epsilon_t^x) / \text{Var}(\epsilon_t^x)$ 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 $\text{Corr}(\epsilon_t^r, \epsilon_t^x) = -0.96$ for the dividend yield as predictor, invalidating standard estimates and tests. For the variable $\Delta\hat{\sigma}_t$, this bias is less of a concern since it is not a price scaled variable. For our one period forecasts, we estimate $\text{Corr}(\Delta\hat{\sigma}_t, \Delta\hat{\sigma}_{t-1}) = 0.006$ and $\text{Corr}(\Delta\hat{\sigma}_t, \epsilon_t) = 0.032$, which is too small to bias statistical inference.

Table XI
Effect of X-12-ARIMA Filter

We simulate 300 time-series of 200 quarterly log consumption growth rates generated from the Markov model estimated in Table I. We then perturb every fourth quarter data point by +5% and every first quarter data point by -5%, which approximately generates the seasonality observed in the discontinued series for seasonally unadjusted quarterly consumption data. We transform the seasonally perturbed series into consumption levels and apply the X-12-ARIMA filter to get seasonally adjusted consumption data. We lastly estimate the four-state Markov model of Section I on both the original and the seasonally adjusted data.

Panel A: Summary Statistics for estimates of undisturbed and X-12 data								
	μ_l	μ_h	σ_l	σ_h	p_μ^{ll}	p_μ^{hh}	p_σ^{ll}	p_σ^{hh}
Mean	0.36	0.79	0.20	0.49	92.26	93.83	92.39	93.77
Mean X-12	0.37	0.80	0.19	0.47	91.94	93.52	90.36	89.96
Median	0.36	0.79	0.20	0.49	94.11	95.67	94.76	95.95
Median X-12	0.36	0.79	0.18	0.46	93.38	94.90	93.64	93.47
SD	0.08	0.06	0.04	0.06	8.18	7.98	9.07	6.91
SD X-12	0.08	0.05	0.04	0.10	6.38	5.09	12.40	11.70
5th Percent	0.28	0.71	0.15	0.41	77.67	85.49	78.57	79.93
5th Percent X-12	0.27	0.72	0.13	0.38	79.36	84.83	66.38	67.92
95th Percent	0.47	0.87	0.27	0.57	98.30	98.91	98.75	99.19
95th Percent X-12	0.48	0.88	0.24	0.57	97.89	98.66	98.60	99.01
Panel B: Median Correlations of beliefs from undisturbed and X-12 data								
	ρ_μ	0.98			ρ_σ	0.90		

Table XII
Wealth-Consumption Ratio

We simulate 300 economies for 100 years at quarterly frequency. In all three panels, the representative agent has an EIS of 1.5 and rate of time preference of 0.995. The coefficient of relative risk aversion (RRA) increases from 10 (Panel A) to 20 (Panel B) and 30 (Panel C). In the first regression of each panel, we regress the log wealth-consumption ratio, pc_t , on the prior probabilities of being in a given state, $\xi_{t+1,t}(i)$, $i = 1, 2, 3$. In the second regression, we regress the log wealth-consumption ratio, pc_t , on the perceived first, $\hat{\mu}_t$, and second moment, $\hat{\sigma}_t$, of consumption growth. We report the average regression coefficient and average R^2 .

Const.	$\xi(1)$	$\xi(2)$	$\xi(3)$	$\hat{\mu}$	$\hat{\sigma}$	R^2
Panel A: RRA=10, EIS=1.5						
5.7172	-0.0056	-0.0074	0.0009			0.9973
5.7115				0.0069	-0.0015	0.9948
Panel B: RRA=20, EIS=1.5						
5.6922	-0.0050	-0.0071	0.0012			0.9957
5.6870				0.0067	-0.0018	0.9933
Panel C: RRA=30, EIS=1.5						
5.6696	-0.0044	-0.0068	0.0016			0.9937
5.6650				0.0064	-0.0021	0.9914

Table XIII
Markov Model of Dividend Growth

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

$$\Delta d_{t+1} = \mu_t^d + \sigma_t^d \epsilon_{t+1} \quad \epsilon_t \sim \mathcal{N}(0, 1)$$

where $\mu_t^d \in \{\mu_l^d, \mu_h^d\}$ and $\sigma_t^d \in \{\sigma_l^d, \sigma_h^d\}$ follow independent Markov processes with transition matrices P^μ and P^σ , respectively. The consumption and dividend process follow the same Markov switching process as reported in Table I. We compute quarterly dividends for the period 1955-2008 using the value-weighted CRSP index with and without distributions. Standard errors are reported in parentheses.

Parameter	Estimate (%)
μ_l^d	-0.4549 (0.1654)
μ_h^d	1.4378 (0.2091)
σ_l^d	1.1747 (0.0838)
σ_h^d	3.4576 (0.4980)

Table XIV
Model Implications

We simulate 300 economies for 100 years at quarterly frequency. In all three panels, the representative agent has an EIS of 1.5 and rate of time preference of 0.995. The coefficient of relative risk aversion (RRA) increases from 10 (Panel A) to 20 (Panel B) and 30 (Panel C). We report the average excess return, $\mathbb{E}[R^e]$, the standard deviation of stock returns, $\sigma[R]$, the average risk-free rate, $\mathbb{E}[R^f]$, and the standard deviation of the risk-free rate, $\sigma[R^f]$. In the last row of each panel, we also report moments of the Markov switching model without learning where the agent knows the state of the economy.

	$\mathbb{E}[R^e]$	$\sigma[R^e]$	$\mathbb{E}[R^f]$	$\sigma[R^f]$
Panel A: RRA=10, EIS=1.5				
Mean	0.0105	0.0599	0.0315	0.0032
25%	0.0058	0.0563	0.0307	0.0030
75%	0.0150	0.0634	0.0323	0.0034
No Learning	0.0098	0.0719	0.0336	0.0027
Panel B: RRA=20, EIS=1.5				
Mean	0.0235	0.0560	0.0274	0.0053
25%	0.0182	0.0523	0.0264	0.0049
75%	0.0285	0.0594	0.0286	0.0058
No Learning	0.0191	0.0700	0.0329	0.0026
Panel C: RRA=30, EIS=1.5				
Mean	0.0361	0.0530	0.0241	0.0071
25%	0.0309	0.0492	0.0225	0.0064
75%	0.0413	0.0566	0.0255	0.0079
No Learning	0.0273	0.0685	0.0322	0.0026