An Intertemporal CAPM with Stochastic Volatility

John Y. Campbell, Stefano Giglio, Christopher Polk, and Robert Turley[1]

First draft: October 2011
This Version: March 2012

[1]Campbell: Department of Economics, Littauer Center, Harvard University, Cambridge MA 02138, and NBER. Email john_campbell@harvard.edu. Phone 617-496-6448. Giglio: Booth School of Business, University of Chicago, 5807 S. Woodlawn Ave, Chicago IL 60637. Email stefano.giglio@chicagobooth.edu. Polk: Department of Finance, London School of Economics, London WC2A 2AE, UK. Email c.polk@lse.ac.uk. Turley: Baker Library 220D, Harvard Business School, Boston MA 02163. Email turley@fas.harvard.edu.
We are grateful to Bjorn Eraker, Ian Martin, Sydney Ludvigson, Monika Piazzesi, Tuomo Vuolteenaho, and seminar participants at INSEAD, University of Paris Dauphine, University of Edinburgh, and the HBS Finance Unit research retreat for comments. We thank Josh Coval, Ken French, Mila Getmansky Sherman, and Tyler Shumway for providing data used in the analysis.

Abstract

This paper extends the approximate closed-form intertemporal capital asset pricing model of Campbell (1993) to allow for stochastic volatility. The return on the aggregate stock market is modeled as one element of a vector autoregressive (VAR) system, and the volatility of all shocks to the VAR is another element of the system. The paper presents evidence that growth stocks underperform value stocks because they hedge two types of deterioration in investment opportunities: declining expected stock returns, and increasing volatility. Volatility hedging is also relevant for pricing risk-sorted portfolios and non-equity assets such as equity index options and corporate bonds.

JEL classification: G12, N22

1 Introduction

The fundamental insight of intertemporal asset pricing theory is that long-term investors should care just as much about the returns they earn on their invested wealth as about the level of that wealth. In a simple model with a constant rate of return, for example, the sustainable level of consumption is the return on wealth multiplied by the level of wealth, and both terms in this product are equally important. In a more realistic model with time-varying investment opportunities, conservative long-term investors will seek to hold "intertemporal hedges", assets that perform well when investment opportunities deteriorate. Such assets should deliver lower average returns in equilibrium if they are priced from conservative long-term investors' first-order conditions.

Since the seminal work of Merton (1973) on the intertemporal capital asset pricing model (ICAPM), a large empirical literature has explored the relevance of intertemporal considerations for the pricing of financial assets in general, and the cross-sectional pricing of stocks in particular. One strand of this literature uses the approximate accounting identity of Campbell and Shiller (1988a) and the assumption that a representative investor has Epstein-Zin utility (Epstein and Zin 1989) to obtain approximate closed-form solutions for the ICAPM's risk prices (Campbell 1993). These solutions can be implemented empirically if they are combined with vector autoregressive (VAR) estimates of asset return dynamics (Campbell 1996). Campbell and Vuolteenaho (2004), Campbell, Polk, and Vuolteenaho (2010), and Campbell, Giglio, and Polk (2011) use this approach to argue that value stocks outperform growth stocks on average because growth stocks do well when the expected return on the aggregate stock market declines; in other words, growth stocks have low risk premia because they are intertemporal hedges for long-term investors.

A weakness of the papers cited above is that they ignore time-variation in the volatility of stock returns. In general, investment opportunities may deteriorate either because expected stock returns decline or because the volatility of stock returns increases, and it is an empirical question which of these two types of intertemporal risk have a greater effect on asset returns. We address this weakness in this paper by extending the approximate closed-form ICAPM to allow for stochastic volatility. The resulting model explains risk premia in the stock market using three priced risk factors corresponding to three important attributes of aggregate market returns: revisions in expected future cash flows, discount rates, and volatility. An attractive characteristic of the model is that the prices of these three risk factors depend on only one free parameter, the long-horizon investor's coefficient of risk aversion.

Our work can be regarded as complementary to recent research on the "long-run risk model" of asset prices (Bansal and Yaron 2004) which can be traced back to insights in Kandel and Stambaugh (1991). Both the approximate closed-form ICAPM and the long-run risk model start with the first-order conditions of an infinitely lived Epstein-Zin representative investor. As originally stated by Epstein and Zin (1989), these first-order conditions involve both aggregate consumption growth and the return on the market portfolio of aggregate wealth. Campbell (1993) pointed out that the intertemporal budget constraint could be used to substitute out consumption growth, turning the model into a Merton-style ICAPM. Restoy and Weil (1998, 2011) used the same logic to substitute out the market portfolio return, turning the model into a generalized consumption CAPM in the style of Breeden (1979).

Kandel and Stambaugh (1991) were the first researchers to study the implications for asset returns of time-varying first and second moments of consumption growth in a model with a representative Epstein-Zin investor. Specifically, Kandel and Stambaugh (1991) assumed a four-state Markov chain for the expected growth rate and conditional volatility of consumption, and provided closed-form solutions for important asset-pricing moments. In the spirit of Kandel and Stambaugh (1991), Bansal and Yaron (2004) added stochastic volatility to the Restoy-Weil model, and subsequent research on the long-run risk model has increasingly emphasized the importance of stochastic volatility for generating empirically plausible implications from this model (Bansal, Kiku, and Yaron 2012, Beeler and Campbell 2012). In this paper we give the approximate closed-form ICAPM the same capability to handle stochastic volatility that its cousin, the long-run risk model, already possesses.

One might ask whether there is any reason to work with an ICAPM rather than a consumption-based model given that these models are derived from the same set of assumptions. The ICAPM developed in this paper has several advantages. First, it describes risks as they appear to an investor who takes asset prices as given and chooses consumption to satisfy his budget constraint. This is the way risks appear to individual agents in the economy, and it seems important for economists to understand risks in the same way that market participants do rather than relying exclusively on a macroeconomic perspective. Second, the ICAPM allows an empirical analysis based on financial proxies for the aggregate market portfolio rather than on accurate measurement of aggregate consumption. While there are certainly challenges to the accurate measurement of financial wealth, financial time series are generally available on a more timely basis and over longer sample periods than consumption series. Third, the ICAPM in this paper is flexible enough to allow multiple state variables that can be estimated in a VAR system; it does not require low-dimensional calibration of the sort used in the long-run risk literature. Finally, the stochastic volatility process used here governs the volatility of all state variables, including itself. We show that this assumption fits financial data reasonably well, and it guarantees that stochastic volatility would always remain positive in a continuous-time version of the model, a property that does not hold in most current implementations of the long-run risk model.[2]

The closest precursors to our work are unpublished papers by Chen (2003) and Sohn (2010).[3] Both papers explore the effects of stochastic volatility on asset prices in an ICAPM setting but make strong assumptions about the covariance structure of various news terms when deriving their pricing equations. Chen (2003) assumes constant covariances between

[2] Eraker (2008) and Eraker and Shaliastovich (2008) are exceptions.

[3] Eraker and Wang (2011) provides an example of a present-value ICAPM where prices respond endogenously to market volatility shocks.

shocks to the market return (and powers of those shocks) and news about future expected market return variance. Sohn (2010) makes two strong assumptions about asset returns and consumption growth, specifically that all assets have zero covariance with news about future consumption growth volatility and that the conditional contemporaneous correlation between the market return and consumption growth is constant through time. Duffee (2005) presents evidence against the latter assumption. It is in any case unattractive to make assumptions about consumption growth in an ICAPM that does not require accurate measurement of consumption.

Chen estimates a VAR with a GARCH model to allow for time variation in the volatility of return shocks, restricting market volatility to depend only on its past realizations and not those of the other state variables. His empirical analysis has little success in explaining the cross-section of stock returns. Sohn uses a similar but more sophisticated GARCH model for market volatility and tests how well short-run and long-run risk components from the GARCH estimation can explain the returns of various stock portfolios, comparing the results to factors previously shown to be empirically successful. In contrast, our paper incorporates the volatility process directly in the ICAPM, allowing heteroskedasticity to affect and to be predicted by all state variables, and showing how the price of volatility risk is pinned down by the time-series structure of the model along with the investor's coefficient of risk aversion.

A working paper by Bansal, Kiku, Shaliastovich and Yaron (2011), contemporaneous with our own, explores the effects of stochastic volatility in the long-run risk model. Like us, they find stochastic volatility to be an important feature in the time series of equity returns. Their work puts greater emphasis on the implied consumption dynamics while we focus on the cross-sectional pricing implications of exposure to volatility news. More fundamentally, there are differences in the underlying models. They assume that the stochastic process driving volatility is homoskedastic, and in their cross-sectional analysis they impose that changes in the equity risk premium are driven only by the conditional variance of the stock market. The different modeling assumptions help explain our contrasting empirical results; we show that volatility risk is very important in explaining the cross-section of stock returns while they find it explains little of the cross-sectional differences in risk premia.

Stochastic volatility has, of course, been explored in other branches of the finance literature. For example, Chacko and Viceira (2005) and Liu (2007) show how stochastic volatility affects the optimal portfolio choice of long-term investors. Chacko and Viceira argue that movements in volatility are not persistent enough to generate large intertemporal hedging demands. Coval and Shumway (2001), Ang, Hodrick, Xing, and Zhang (2006), and Adrian and Rosenberg (2008) present evidence that shocks to market volatility are priced risk factors in the cross-section of stock returns, but they do not develop any theory to explain the risk prices for these factors. There is also a large literature in financial econometrics describing how to use realized volatility from high-frequency data to estimate stochastic volatility processes (Barndorff-Nielsen and Shephard 2002, Andersen, Bollerslev, Diebold, and Labys 2003). We follow this literature by including a realized volatility measure in our VAR.

The empirical implementation of our model is a success. We find that growth stocks have low average returns because they outperform not only when the expected stock return declines, but also when stock market volatility increases. Thus growth stocks hedge two types of deterioration in investment opportunities, not just one. In the period since 1963 that creates the greatest empirical difficulties for the standard CAPM, we find that the three-beta model explains over 72% of the cross-sectional variation in average returns of 25 portfolios sorted by size and book-to-market ratios. The model is not rejected at the 5% level while the CAPM is strongly rejected. The implied coefficient of relative risk aversion is an economically reasonable 9.13, in contrast to the much larger estimate of 19.90, which we get when we estimate the two-beta CAPM of Campbell and Vuolteenaho (2004) using the same data.[4] This success is due in large part to the inclusion of volatility betas in the specification. In particular, the spread in volatility betas in the cross section generates an annualized spread in average returns of 7.84% compared to a comparable spread of 3.48% and 2.20% for cash-flow and discount-rate betas.

We confirm that our findings are robust by expanding the set of test portfolios beyond the 25 size- and book-to-market-sorted portfolios in two important dimensions. First, we show that our three-beta model not only describes the cross section of characteristics-sorted portfolios but also can explain the average returns on risk-sorted portfolios. We examine risk-sorted portfolios in response to the argument of Daniel and Titman (1997, 2012) and Lewellen, Nagel, and Shanken (2010) that asset-pricing tests using only portfolios sorted by characteristics known to be related to average returns, such as size and value, can be misleading. As tests that include risk-sorted portfolios are unable to reject our intertemporal CAPM with stochastic volatility, we verify that the model's success is not simply due to the low-dimensional factor structure of the 25 size- and book-to-market-sorted portfolios. Specifically, we show that sorts on stocks' pre-formation sensitivity to volatility news generate economically and statistically significant spread in both post-formation volatility beta and average returns in a manner consistent with our model. Interestingly, in the post-1963 period, sorts on past CAPM beta generate little spread in post-formation cash-flow betas, but significant spread in post-formation volatility betas. Since, in the three-beta model, covariation with aggregate volatility news has a negative premium, the three-beta model also explains why stocks with high past CAPM betas have offered relatively little extra return in the post-1963 sample.

Second, we show that our three-beta model can help explain average returns on non-equity portfolios that are exposed to aggregate volatility risk. These portfolios include the S&P 100 index straddle of Coval and Shumway (2001), which is explicitly designed to be highly correlated with aggregate volatility risk, and the risky bond factor of Fama and French (1993), which should be sensitive to changes in aggregate volatility since risky corporate debt is short the option to default. Consistent with this intuition, we find that compared to the volatility beta of a value-minus-growth bet, the risky bond factor's volatility beta is of the same order of magnitude while the straddle's volatility beta is more than 2.7 times larger

[4]The risk aversion estimate reported in Campbell and Vuolteenaho's (2004) paper is 28.75.

in absolute magnitude. These volatility betas are of the right sign to explain the abnormal CAPM returns of the option and bond portfolios. Approximately 25% of the average straddle return can be attributed to its three ICAPM betas, based purely on model estimates from the cross section of equity returns. Additionally, when we price the *joint* cross-section of equity, bond, and straddle returns our intertemporal CAPM with stochastic volatility is not rejected at the 5-percent level while the CAPM is strongly rejected.

The organization of our paper is as follows. Section 2 lays out the approximate closed-form ICAPM and shows how to extend it to incorporate stochastic volatility. Section 3 presents data, econometrics, and VAR estimates of the dynamic process for stock returns and realized volatility. Section 4 turns to cross-sectional asset pricing and estimates a representative investor's preference parameters to fit the cross-section, taking the dynamics of stock returns as given. Section 5 concludes.

2 An Intertemporal Model with Stochastic Volatility

2.1 Asset Pricing with Time Varying Risk

Preferences

We begin by assuming a representative agent with Epstein-Zin preferences. We write the value function as

$$V_t = \left[(1-\delta)\, C_t^{\frac{1-\gamma}{\theta}} + \delta\left(\mathrm{E}_t\left[V_{t+1}^{1-\gamma}\right]\right)^{1/\theta}\right]^{\frac{\theta}{1-\gamma}}, \tag{1}$$

where C_t is consumption and the preference parameters are the discount factor δ, risk aversion γ, and the elasticity of intertemporal substitution ψ. For convenience, we define $\theta = (1-\gamma)/(1-1/\psi)$.

The corresponding stochastic discount factor (SDF) can be written as

$$M_{t+1} = \left(\delta\left(\frac{C_t}{C_{t+1}}\right)^{1/\psi}\right)^{\theta}\left(\frac{W_t - C_t}{W_{t+1}}\right)^{1-\theta}, \tag{2}$$

where W_t is the market value of the consumption stream owned by the agent, including current consumption C_t.[5] The log return on wealth is $r_{t+1} = \ln\left(W_{t+1}/\left(W_t - C_t\right)\right)$, the log value of wealth tomorrow divided by reinvested wealth today. The log SDF is therefore

$$m_{t+1} = \theta \ln \delta - \frac{\theta}{\psi}\Delta c_{t+1} + (\theta - 1)\, r_{t+1}. \tag{3}$$

[5] This notational convention is not consistent in the literature. Some authors exclude current consumption from the definition of current wealth.

A convenient identity

The gross return to wealth can be written

$$1 + R_{t+1} = \frac{W_{t+1}}{W_t - C_t} = \left(\frac{C_t}{W_t - C_t}\right)\left(\frac{C_{t+1}}{C_t}\right)\left(\frac{W_{t+1}}{C_{t+1}}\right), \tag{4}$$

expressing it as the product of the current consumption payout, the growth in consumption, and the future price of a unit of consumption.

We find it convenient to work in logs. We define the log value of reinvested wealth per unit of consumption as $z_t = \ln\left(\left(W_t - C_t\right)/C_t\right)$, and the future value of a consumption claim as $h_{t+1} = \ln\left(W_{t+1}/C_{t+1}\right)$, so that the log return is:

$$r_{t+1} = -z_t + \Delta c_{t+1} + h_{t+1}. \tag{5}$$

Heuristically, the return on wealth is negatively related to the current value of reinvested wealth and positively related to consumption growth and the future value of wealth. The last term in equation (5) will capture the effects of intertemporal hedging on asset prices, hence the choice of the notation h_{t+1} for this term.

The ICAPM

We assume that asset returns are jointly conditionally lognormal, but we allow changing conditional volatility so we are careful to write second moments with time subscripts to indicate that they can vary over time. Under this standard assumption, the expected return on any asset must satisfy

$$0 = \ln \mathrm{E}_t \exp\{m_{t+1} + r_{i,t+1}\} = \mathrm{E}_t\left[m_{t+1} + r_{i,t+1}\right] + \frac{1}{2}\mathrm{Var}_t\left[m_{t+1} + r_{i,t+1}\right], \tag{6}$$

and the risk premium on any asset is given by

$$\mathrm{E}_t r_{i,t+1} - r_{f,t} + \frac{1}{2}\mathrm{Var}_t r_{t+1} = -\mathrm{Cov}_t\left[m_{t+1}, r_{i,t+1}\right]. \tag{7}$$

The convenient identity (5) can be used to write the log SDF (3) without reference to consumption growth:

$$m_{t+1} = \theta \ln \delta - \frac{\theta}{\psi} z_t + \frac{\theta}{\psi} h_{t+1} - \gamma r_{t+1}. \tag{8}$$

Since the first two terms in (5) are known at time t, only the latter two terms appear in the conditional covariance in (7). We obtain an ICAPM pricing equation that relates the risk premium on any asset to the asset's covariance with the wealth return and with shocks to future consumption claim values:

$$\mathrm{E}_t r_{i,t+1} - r_{f,t} + \frac{1}{2}\mathrm{Var}_t r_{t+1} = \gamma \mathrm{Cov}_t\left[r_{i,t+1}, r_{t+1}\right] - \frac{\theta}{\psi}\mathrm{Cov}_t\left[r_{i,t+1}, h_{t+1}\right] \tag{9}$$

Return and Risk Shocks in the ICAPM

To better understand the intertemporal hedging component h_{t+1}, we proceed in two steps. First, we approximate the relationship of h_{t+1} and z_{t+1} by taking a loglinear approximation about $\bar{z}$:

$$h_{t+1} \approx \kappa + \rho z_{t+1} \tag{10}$$

where the loglinearization parameter $\rho = \exp(\bar{z})/(1+\exp(\bar{z})) \approx 1 - C/W$.

Second, we apply the general pricing equation (6) to the wealth portfolio itself (setting $r_{i,t+1} = r_{t+1}$), and use the convenient identity (5) to substitute out consumption growth from this expression. Rearranging, we can write the variable z_t as

$$z_t = \psi \ln \delta + (\psi - 1)\mathrm{E}_t r_{t+1} + \mathrm{E}_t h_{t+1} + \frac{\psi}{\theta}\frac{1}{2}\mathrm{Var}_t\left[m_{t+1} + r_{t+1}\right]. \tag{11}$$

Third, we combine these expressions to obtain the innovation in h_{t+1}:

$$\begin{aligned} h_{t+1} - \mathrm{E}_t h_{t+1} &= \rho(z_{t+1} - \mathrm{E}_t z_{t+1}) \\ &= (\mathrm{E}_{t+1} - \mathrm{E}_t)\rho\left((\psi - 1)r_{t+2} + h_{t+2} + \frac{\psi}{\theta}\frac{1}{2}\mathrm{Var}_{t+1}\left[m_{t+2} + r_{t+2}\right]\right). \end{aligned} \tag{12}$$

Solving forward to an infinite horizon,

$$\begin{aligned} h_{t+1} - \mathrm{E}_t h_{t+1} &= (\psi - 1)(\mathrm{E}_{t+1} - \mathrm{E}_t)\sum_{j=1}^{\infty}\rho^j r_{t+1+j} \\ &\quad + \frac{1}{2}\frac{\psi}{\theta}(\mathrm{E}_{t+1} - \mathrm{E}_t)\sum_{j=1}^{\infty}\rho^j \mathrm{Var}_{t+j}\left[m_{t+1+j} + r_{t+1+j}\right] \\ &= (\psi - 1)N_{DR,t+1} + \frac{1}{2}\frac{\psi}{\theta}N_{RISK,t+1}. \end{aligned} \tag{13}$$

The second equality follows Campbell and Vuolteenaho (2004) and uses the notation N_{DR} ("news about discount rates") for revisions in expected future returns. In a similar spirit we write revisions in expectations of future risk (the variance of the future log return plus the log stochastic discount factor) as N_{RISK}.

Finally, we substitute back into the intertemporal model (9):

$$\begin{aligned} &\mathrm{E}_t r_{i,t+1} - r_{f,t} + \frac{1}{2}\mathrm{Var}_t r_{i,t+1} \\ = \;& \gamma \mathrm{Cov}_t\left[r_{i,t+1}, r_{t+1}\right] + (\gamma - 1)\,\mathrm{Cov}_t\left[r_{i,t+1}, N_{DR,t+1}\right] - \frac{1}{2}\mathrm{Cov}_t\left[r_{i,t+1}, N_{RISK,t+1}\right] \\ = \;& \gamma \mathrm{Cov}_t\left[r_{i,t+1}, N_{CF,t+1}\right] + \mathrm{Cov}_t\left[r_{i,t+1}, -N_{DR,t+1}\right] - \frac{1}{2}\mathrm{Cov}_t\left[r_{i,t+1}, N_{RISK,t+1}\right]. \end{aligned} \tag{14}$$

The first equality expresses the risk premium as risk aversion γ times covariance with the current market return, plus $(\gamma - 1)$ times covariance with news about future market returns, minus one half covariance with risk. This is an extension of the ICAPM as written by Campbell (1993), with no reference to consumption or the elasticity of intertemporal substitution ψ.[6] When the investor's risk aversion is greater than 1, assets which hedge aggregate discount rates ($\mathrm{Cov}_t\left[r_{i,t+1}, N_{DR,t+1}\right] < 0$) or aggregate risk ($\mathrm{Cov}_t\left[r_{i,t+1}, N_{RISK,t+1}\right] > 0$) have lower expected returns, all else equal.

The second equality rewrites the model, following Campbell and Vuolteenaho (2004), by breaking the market return into cash-flow news and discount-rate news. Cash-flow news N_{CF} is defined by $N_{CF} = r_{t+1} - \mathrm{E}_t r_{t+1} + N_{DR}$. The price of risk for cash-flow news is γ times greater than the price of risk for discount-rate news, hence Campbell and Vuolteenaho call betas with cash-flow news "bad betas" and those with discount-rate news "good betas" since they have lower risk prices in equilibrium. The third term in (14) shows the risk premium associated with exposure to news about future risks and did not appear in Campbell and Vuolteenaho's model, which assumed homoskedasticity. Not surprisingly, the coefficient is negative, indicating that an asset providing positive returns when risk expectations increase will offer a lower return on average.

2.2 From Risk to Volatility

The risk shocks defined in the previous subsection are shocks to the conditional volatility of returns plus the stochastic discount factor, that is, the conditional volatility of risk-neutralized returns. We now make additional assumptions to derive a model in which this conditional volatility is proportional to the conditional volatility of returns themselves.

Suppose the economy is described by a first-order VAR

$$\mathbf{x}_{t+1} = \bar{\mathbf{x}} + \mathbf{\Gamma}\left(\mathbf{x}_t - \bar{\mathbf{x}}\right) + \sigma_t \mathbf{u}_{t+1}, \tag{15}$$

where $\mathbf{x}_{t+1}$ is an $n \times 1$ vector of state variables that has r_{t+1} as its first element, σ^2_{t+1} as its second element, and $n-2$ other variables that help to predict the first and second moments of aggregate returns. $\boldsymbol{\phi}$ and $\mathbf{\Gamma}$ are an $n \times 1$ vector and an $n \times n$ matrix of constant parameters, and $\mathbf{u}_{t+1}$ is a vector of shocks to the state variables. The key assumption here is that a scalar random variable, σ^2_t, governs time-variation in the variance of all shocks to this system. Both market returns and state variables, including volatility itself, have innovations whose variances move in proportion to one another.

[6] Campbell (1993) briefly considers the heteroskedastic case, noting that when $\gamma = 1$, $\mathrm{Var}_t\left[m_{t+1} + r_{t+1}\right]$ is a constant. This implies that N_{RISK} does not vary over time so the stochastic volatility term disappears. Campbell claims that the stochastic volatility term also disappears when $\psi = 1$, but this is incorrect. When limits are taken correctly, N_{RISK} does not depend on ψ (except indirectly through the loglinearization parameter, ρ).

Given this structure, news about discount rates can be written as

$$\begin{aligned} N_{DR,t+1} &= (E_{t+1}-E_t)\sum_{j=1}^{\infty}\rho^j r_{t+1+j} \\ &= \mathbf{e}_1'\sum_{j=1}^{\infty}\rho^j\mathbf{\Gamma}^j\sigma_t\mathbf{u}_{t+1} \\ &= \mathbf{e}_1'\rho\mathbf{\Gamma}\left(\mathbf{I}-\rho\mathbf{\Gamma}\right)^{-1}\sigma_t\mathbf{u}_{t+1} \end{aligned} \tag{16}$$

Furthermore, our log-linear model will make the log SDF, m_{t+1}, a linear function of the state variables. Since all shocks to the SDF are then proportional to σ_t, $\mathrm{Var}_t\left[m_{t+1}+r_{t+1}\right]\propto\sigma_t^2$. As a result, the conditional variance, $\mathrm{Var}_t\left[\left(m_{t+1}+r_{t+1}\right)/\sigma_t\right]=\omega_t$, will be a constant that does not depend on the state variables. Without knowing the parameters of the utility function, we can write $\mathrm{Var}_t\left[m_{t+1}+r_{t+1}\right]=\omega\sigma_t^2$ so that the news about risk, N_{RISK}, is proportional to news about market return variance, N_V.

$$\begin{aligned} N_{RISK,t+1} &= (E_{t+1}-E_t)\sum_{j=1}^{\infty}\rho^j\mathrm{Var}_{t+j}\left[r_{t+1+j}+m_{t+1+j}\right] \\ &= (E_{t+1}-E_t)\sum_{j=1}^{\infty}\rho^j\left(\omega\sigma_{t+j}^2\right) \\ &= \omega\rho\mathbf{e}_2'\sum_{j=0}^{\infty}\rho^j\mathbf{\Gamma}^j\sigma_t\mathbf{u}_{t+1} \\ &= \omega\rho\mathbf{e}_2'\left(\mathbf{I}-\rho\mathbf{\Gamma}\right)^{-1}\sigma_t\mathbf{u}_{t+1}=\omega N_{V,t+1}. \end{aligned} \tag{17}$$

Substituting (17) into (14), we obtain an empirically-testable intertemporal CAPM with stochastic volatility:

$$\begin{aligned} &\mathrm{E}_t r_{i,t+1}-r_{f,t}+\frac{1}{2}\mathrm{Var}_t r_{i,t+1} \\ =\ &\gamma\mathrm{Cov}_t\left[r_{i,t+1},N_{CF,t+1}\right]+\mathrm{Cov}_t\left[r_{i,t+1},-N_{DR,t+1}\right]-\frac{1}{2}\omega\mathrm{Cov}_t\left[r_{i,t+1},N_{V,t+1}\right], \end{aligned} \tag{18}$$

where covariances with news about three key attributes of the market portfolio (cash flows, discount rates, and volatility) describe the cross section of average returns.

The parameter ω is a nonlinear function of the coefficient of relative risk aversion γ, as well as the VAR parameters and the loglinearization coefficient ρ, but it does not depend on the elasticity of intertemporal substitution ψ except indirectly through the influence of ψ on ρ. In the appendix, we show that ω solves:

$$\omega\sigma_t^2=(1-\gamma)^2\mathrm{Var}_t\left[N_{CF_{t+1}}\right]+\omega(1-\gamma)\mathrm{Cov}_t\left[N_{CF_{t+1},}N_{V_{t+1},}\right]+\omega^2\frac{1}{4}\mathrm{Var}_t\left[N_{V_{t+1}}\right]. \tag{19}$$

We can see two main channels through which γ affects ω. First, a higher risk aversion—given the underlying volatilities of all shocks—implies a more volatile stochastic discount factor m, and therefore a higher RISK. This effect is proportional to $(1-\gamma)^2$, so it increases rapidly with γ. Second, there is a feedback effect on RISK through future risk: ω appears on the right-hand side of the equation as well. Given that in our estimation we find $\text{Cov}_t\left[N_{CF_{t+1}}, N_{V_{t+1},}\right] < 0$, this second effect makes ω increase even faster with γ.[7]

This equation can also be written directly in terms of the VAR parameters. If we define x_{CF} and x_V as the error-to-news vectors such that

$$\frac{1}{\sigma_t} N_{CF,t+1} = x_{CF} u_{t+1} = \left(e_1' + e_1' \rho\Gamma(I-\rho\Gamma)^{-1}\right) u_{t+1} \tag{20}$$

$$\frac{1}{\sigma_t} N_{V,t+1} = x_V u_{t+1} = \left(e_2' \rho(I-\rho\Gamma)^{-1}\right) u_{t+1} \tag{21}$$

and define the covariance matrix of the residuals (scaled to eliminate stochastic volatility) as $\mathbf{\Sigma} = \text{Var}[\mathbf{u}_{t+1}]$, then ω solves

$$0 = \omega^2 \frac{1}{4} x_V \mathbf{\Sigma} x_V' - \omega \left(1 - (1-\gamma)\, x_{CF} \mathbf{\Sigma} x_V'\right) + (1-\gamma)^2\, x_{CF} \mathbf{\Sigma} x_{CF}' \tag{22}$$

This quadratic equation for ω has two solutions. This result is an artifact of our linear approximation of the Euler Equation, and the appendix shows that one of the solutions can be disregarded. This false solution is easily identified by its implication that ω becomes infinite as volatility shocks become small. The correct solution is

$$\omega = \frac{1-(1-\gamma)x_{CF}\Sigma x_V' - \sqrt{(1-(1-\gamma)x_{CF}\Sigma x_V')^2 - (1-\gamma)^2 (x_V \Sigma x_V')(x_{CF}\Sigma x_{CF}')}}{\frac{1}{2} x_V \Sigma x_V'} \tag{23}$$

There is an additional disadvantage to the quadratic expression arising from our loglinearization. In the case where risk aversion, volatility shocks and cash flow shocks are large enough, as measured by the product $(1-\gamma)^2(x_V\Sigma x_V')(x_{CF}\Sigma x_{CF}')$. equation (22) may deliver a complex rather than a real value for ω. While the conditional variance $\text{Var}_t[m_{t+1}+r_{t+1}]$ from which we define ω will be both real and finite, the loglinear approximation may not allow for a real solution in an economically important region of the parameter space. Given our VAR estimates of the variance and covariance terms, we find equation (22) yields a real solution as γ ranges from zero to 6.39.

To allow for larger values in our risk aversion parameter, we consider an alternative approximation. If we linearize the right hand side of equation (19) around $\omega = 0$ we can

[7] Bansal, Kiku, Shaliastovich and Yaron (2011) derive a similar expression. The equivalent expression for ω in their case reduces to $(1-\gamma)^2$ as they impose that the volatility process is homoskedastic and the conditional equity premium is driven solely by the stochastic volatility.

approximate $\text{Var}_t[m_{t+1} + r_{t+1}]$ as a linear, rather than quadratic, function of ω. We then have

$$\omega \approx \frac{(1-\gamma)^2 (x_{CF} \Sigma x'_{CF})}{1 - (1-\gamma)(x_{CF} \Sigma x'_V)} \tag{24}$$

which is now defined for all $\gamma > 0$. Figure 1 plots ω as a function of γ using both the solution in equation (23) and the approximation in (24) for values of γ up to 15.

By construction, they will yield similar solutions for values of γ close to one, where ω gets close to 0 and volatility news becomes less and less important. In other words, it is easy to show that our linearization preserves the property of the true model that as $\gamma \to 1$, $\omega \to 0$ and

$$\text{Var}_t[m_{t+1} + r_{t+1}] \to (1-\gamma)^2 \text{Var}_t[N_{CF}]$$

As risk aversion increases, we find that this approximate value for ω continues to resemble the exact solution of the quadratic equation (22) in the region where a real solution exists. We have also used numerical methods, similar to those proposed by Tauchen and Hussey (1991), to solve the model and validate our estimates of ω for a range of values for γ that include the region where the quadratic equation does not have a real solution.

2.3 Implications for consumption growth

Following Campbell (1993), in this paper we substitute consumption out of the pricing equations using the intertemporal budget constraint. However the model does have interesting implications for the implied consumption process. From equations (5) and (13), we can derive the expression:

$$\Delta c_{t+1} - E_t \Delta c_{t+1} = (r_{t+1} - E_t r_{t+1}) - (\psi - 1) N_{DR,t+1} - (\psi - 1) \frac{1}{2} \frac{\omega}{1-\gamma} N_{V,t+1}.$$

The first two components of the equation for consumption growth are the same as in the homoskedastic case. An unexpectedly high return of the wealth portfolio has a one-for-one effect on consumption. An increase in expected future returns increases today's consumption if $\psi < 1$, as the low elasticity of intertemporal substitution induces the representative investor to consume today (the income effect dominates). If $\psi > 1$, instead, the same increase induces the agent to reduce consumption to better exploit the improved investment opportunities (the substitution effect dominates).

The introduction of time-varying conditional volatility adds an additional term to the equation describing consumption growth. News about high future risk is news about a deterioration of future investment opportunities, which is bad news for a risk-averse investor ($\gamma > 1$). When $\psi < 1$, the representative agent will reduce consumption and save to ensure adequate future consumption. An investor with high elasticity of intertemporal substitution, on the other hand, will increase current consumption and reduce the amount of wealth exposed to the future (worse) investment opportunities.

Using estimates of the news terms from our VAR model (described in the next section), we can explore the implications of the model for consumption growth. Since, as shown in the previous subsection, the three shocks that drive innovations in consumption growth ($r_{t+1} - E_t r_{t+1}$, $N_{DR,t+1}$, $N_{V,t+1}$) can all be expressed as functions of the shock $\sigma_t u_{t+1}$, it follows that we can write the conditional variance of consumption growth as:

$$\text{Var}_t(\Delta c_{t+1}) = A(\gamma, \psi)\sigma_t^2,$$

where $A(\gamma, \psi)$ is a function of the variances and covariances of the scaled residuals $\mathbf{u}_{t+1}$.

Figure 2 plots the coefficient $A(\gamma, \psi)$ for different values of γ and ψ for the homoskedastic case (left panel), and for the heteroskedastic case (right panel) using the linear approximation for ω described in Section 2.2. In each panel, we plot $A(\gamma, \psi)$ as γ varies between 0 and 15, for different values of ψ. Each line corresponds to a different ψ between 0.5 and 1.5.

As expected, in the homoskedastic case (left panel), the variance of consumption growth does not depend on γ but only on ψ. It is rising in ψ because our VAR estimates imply that the return on wealth is negatively correlated with news about future expected returns $N_{DR,t+1}$, that is, wealth returns are mean-reverting. This confirms results reported in Campbell (1996). Once we add stochastic volatility (right panel), as γ increases the volatility of consumption growth increases for all values of ψ as long as $\psi \neq 1$. To understand why this is the case, notice in equation (24) that since ω grows with γ faster than $(1-\gamma)^2$, the term $\frac{\omega}{1-\gamma}$ is increasing in γ in absolute value. Therefore, the larger γ, the more the variance of N_V gets amplified into a higher variance of consumption innovations.

Note also that for $\psi < 1$ and for high enough γ (i.e. in the bottom-right section of the right panel), the volatility of consumption innovations is *higher* for *lower* values of ψ. When risk aversion is high, innovations in consumption are dominated by news about future risk. Agents with very low or very high elasticity of intertemporal substitution, i.e. with ψ far from 1, will tend to adjust their consumption strongly (in different directions) to volatility news. Therefore, it is possible for individuals with *lower* elasticity of intertemporal substitution to end up with a *more volatile* process for consumption innovations, due to their strong reaction to volatility news.

3 Data and Econometrics

3.1 State variables and volatility estimation

Our full VAR specification of the vector $\mathbf{x}_{t+1}$ includes six state variables, five of which are the same as in Campbell, Giglio and Polk (2011). To those five variables, we add an estimate of conditional volatility. The data are all quarterly, from 1926:2 to 2010:4.

The first variable in the VAR is the log real return on the market, r_M, the difference between the log return on the Center for Research in Securities Prices (CRSP) value-weighted stock index and the log return on the Consumer Price Index.

The second variable is expected market variance ($EVAR$). This variable is meant to capture the volatility of market returns, σ_t, conditional on information available at time t, so that innovations to this variable can be mapped to the N_V term described above. To construct $EVAR_t$, we proceed as follows. We first construct a series of within-quarter *realized* variance of daily returns for each time t, $RVAR_t$. We then run a regression of $RVAR_{t+1}$ on lagged realized variance ($RVAR_t$) as well as the other five state variables at time t. This regression then generates a series of predicted values for $RVAR$ at each time $t+1$, that depend on information available at time t: $\widehat{RVAR}_{t+1}$. Finally, we define our expected variance at time t to be exactly this predicted value at $t+1$:

$$EVAR_t \equiv \widehat{RVAR}_{t+1}.$$

The third variable is the price-earnings ratio (PE) from Shiller (2000), constructed as the price of the S&P 500 index divided by a ten-year trailing moving average of aggregate earnings of companies in the S&P 500 index. Following Graham and Dodd (1934), Campbell and Shiller (1988b, 1998) advocate averaging earnings over several years to avoid temporary spikes in the price-earnings ratio caused by cyclical declines in earnings. We avoid any interpolation of earnings as well as lag the moving average by one quarter in order to ensure that all components of the time-t price-earnings ratio are contemporaneously observable by time t. The ratio is log transformed.

Fourth, the term yield spread (TY) is obtained from Global Financial Data. We compute the TY series as the difference between the log yield on the 10-Year US Constant Maturity Bond (IGUSA10D) and the log yield on the 3-Month US Treasury Bill (ITUSA3D).

Fifth, the small-stock value spread (VS) is constructed from data on the six "elementary" equity portfolios also obtained from Professor French's website. These elementary portfolios, which are constructed at the end of each June, are the intersections of two portfolios formed on size (market equity, ME) and three portfolios formed on the ratio of book equity to market equity (BE/ME). The size breakpoint for year t is the median NYSE market equity at the end of June of year t. BE/ME for June of year t is the book equity for the last fiscal year end in $t-1$ divided by ME for December of $t-1$. The BE/ME breakpoints are the 30th and 70th NYSE percentiles.

At the end of June of year t, we construct the small-stock value spread as the difference between the $\ln(BE/ME)$ of the small high-book-to-market portfolio and the $\ln(BE/ME)$ of the small low-book-to-market portfolio, where BE and ME are measured at the end of December of year $t-1$. For months from July to May, the small-stock value spread is constructed by adding the cumulative log return (from the previous June) on the small low-book-to-market portfolio to, and subtracting the cumulative log return on the small high-book-to-market portfolio from, the end-of-June small-stock value spread. The construction of this series follows Campbell and Vuolteenaho (2004) closely.

The sixth variable in our VAR is the default spread (DEF), defined as the difference between the log yield on Moody's BAA and AAA bonds. The series is obtained from the Federal Reserve Bank of St. Louis. Campbell, Giglio and Polk (2011) add the default spread to the Campbell and Vuolteenaho (2004) VAR specification in part because that variable is known to track time-series variation in expected real returns on the market portfolio (Fama and French, 1989), but mostly because shocks to the default spread should to some degree reflect news about aggregate default probabilities. Of course, news about aggregate default probabilities should in turn reflect news about the market's future cash flows.

In order for the regression model that generates $EVAR_t$ to be consistent with a reasonable data-generating process for market variance, we deviate from standard OLS in two ways. First, given that we explicitly consider heteroskedasticity of the innovations to our variables, we estimate this regression using Weighted Least Squares (WLS), where the weight of each observation pair $(RVAR_{t+1}, \mathbf{x}_t)$ is initially based on the time-t value of $(RVAR)^{-1}$. However, to ensure that the ratio of weights across observations is not extreme, we shrink these initial weights towards equal weights. In particular, we set our shrinkage factor large enough so that the ratio of the largest observation weight to the smallest observation weight is always bounded by the corresponding ratio observed for the VIX index.[8]

Second, we deviate from OLS by constraining the regression coefficients to produce fitted values (i.e. expected market return variance) that lie between reasonable ex-ante bounds. To be consistent with our constraints on the weights in WLS, we again use the observed range on the VIX index to inform our priors. Both the constraint on observation weights and the constraint on the regression's fitted values bind in the sample we study.

The results of the first stage regression generating the state variable $EVAR_t$ are reported in Table 1. Perhaps not surprisingly, past realized variance strongly predicts future realized variance. In addition, Table 1 documents that an increase in either PE or DEF predicts higher future realized volatility. Both of these results are statistically significant. The result that higher PE predicts *higher* $RVAR$ might seem surprising at first, but one has to remember that the coefficient indicates the effect of a change in PE holding constant the other variables, including the return on the market; therefore, it captures the effect of a decrease in earnings on future volatility. The R^2 of this regression is just over 10%. The low R^2 masks the fact that the fit is indeed quite good, as we can see from Figure 3, in which $RVAR$ and $EVAR$ are plotted together. The R^2 is heavily influenced by the occasional spikes in realized variance, which the simple linear model we use is not able to capture. Indeed, our WLS approach downweights the importance of those spikes in the estimation procedure.

[8]According to the CBOE website, the VIX reached a minimum of 9.31% annualized volatility on 12/22/1993 and a maximum of 80.86 annualized volatility on 11/20/2008.

The internet appendix to this paper (Campbell, Giglio, Polk, and Turley 2012) reports descriptive statistics for these variables for the full sample, the early sample, and the modern sample. Consistent with Campbell, Giglio and Polk (2011), we document high correlation between DEF and both PE and VS. The table also documents the high persistence of both $RVAR$ and $EVAR$ (autocorrelations of 0.525 and 0.757 respectively) and the high correlation between these variance measures and the default spread. Perhaps the most notable difference between the two subsamples is the correlation between PE and $EVAR$. In the early sample, this correlation is strongly negative, with a value of -0.573. This strong negative correlation reflects the high volatility that occurred during the Great Depression when prices were relatively low. In the modern sample, the correlation is much closer to zero, -0.047. This estimate reflects a mix of episodes with high volatility and high stock prices, such as the technology boom of the late 1990s, and episodes with high volatility and low stock prices, such as the recession of the early 1980s.

3.2 Estimation of the VAR and the news terms

Following Campbell (1993), we estimate a first-order VAR as in equation (15), where $\mathbf{x}_{t+1}$ is a 6×1 vector of state variables ordered as follows:

$$\mathbf{x}_{t+1} = [r_{M,t+1}\ EVAR_{t+1}\ PE_{t+1}\ TY_{t+1}\ DEF_{t+1}\ VS_{t+1}]$$

so that the real market return $r_{M,t+1}$ is the first element and $EVAR$ is the second element. $\bar{\mathbf{x}}$ is a 6×1 vector of the means of the variables, and $\mathbf{\Gamma}$ is a 6×6 matrix of constant parameters. Finally, $\sigma_t \mathbf{u}_{t+1}$ is a 6×1 vector of innovations, with the conditional variance-covariance matrix of $\mathbf{u}_{t+1}$ a constant:

$$\mathbf{\Sigma} = \mathrm{Var}(\mathbf{u}_{t+1})$$

so that the parameter σ_t^2 scales the entire variance-covariance matrix of the vector of innovations.

The first-stage regression forecasting realized market return variance described in the previous section generates the variable $EVAR$. The theory in Section 2 assumes that σ_t^2, proxied for by $EVAR$, scales the variance-covariance matrix of state variable shocks. Thus, as in the first stage, we estimate the second-stage VAR using WLS, where the weight of each observation pair ($\mathbf{x}_{t+1}$, $\mathbf{x}_t$) is initially based on $(EVAR_t)^{-1}$. We continue to constrain both the weights across observations and the fitted values of the regression forecasting $EVAR$ to be consistent with the historical properties of the VIX index.

Table 2 Panel A presents the results of the VAR estimation for the full sample (1926:2 to 2010:4). We report both Newey-West standard errors, estimated with a lag length of four quarters, and bootstrap standard errors for the parameter estimates of the VAR. The bootstrap standard errors for our second-stage regression allow us to take into account the uncertainty generated by forecasting variance in the first stage. Consistent with previous research, we find that PE and DEF negatively predict future returns, though DEF is only marginally significant. The value spread, which is highly correlated with both EVAR and the default spread, has a positive but not statistically significant effect on future returns. In our specification, a higher conditional variance, $EVAR$, is associated with higher future returns, though the effect is not statistically significant. Indeed, once the uncertainty generated by the first stage is taken into account, no variable is statistically significant. As for the other novel aspects of the transition matrix, both high PE and high DEF predict higher future conditional variance of returns.

Panel B of Table 2 reports the sample correlation and autocorrelation matrices of both the unscaled residuals $\sigma_t \mathbf{u}_{t+1}$ and the scaled residuals $\mathbf{u}_{t+1}$. The correlation matrices report standard deviations on the diagonals. There are a couple of aspects of these results to note. For one thing, a comparison of the standard deviations of the unscaled and scaled residuals provides a rough indication of the effectiveness of our empirical solution to the heteroskedasticity of the VAR. In general, the standard deviations of the scaled residuals are several times larger than their unscaled counterparts. Our approach implies that the scaled return residuals should have unit standard deviation. Our implementation results in a sample standard deviation smaller than this at 0.552.

Additionally, a comparison of the unscaled and scaled autocorrelation matrices reveals that much of the sample autocorrelation in the unscaled residuals is eliminated by our WLS approach. For example, the unscaled residuals in the regression forecasting the log real return have an autocorrelation of -0.135. The corresponding autocorrelation of the scaled return residuals is essentially zero, -0.003. Similarly, the autocorrelation in $EVAR$ is reduced from -0.088 to -0.002. Though the scaled residuals in the PE and DEF regression still display significant negative autocorrelation, the unscaled residuals are much more negatively autocorrelated.

Table 3 reports the coefficients of a regression of the squared unscaled residuals $\sigma_t u_{t+1}$ of each VAR equation on a constant and $EVAR$. These results are consistent with our assumption that $EVAR$ captures the conditional volatility of market returns (the coefficient of $EVAR$ for the squared residuals of r_M is 0.649 and not significantly different from one, and the intercept is not significantly different from zero). The fact that $EVAR$ significantly predicts with a positive sign all the squared errors of the VAR supports our underlying assumption that one parameter (σ_t^2) drives the volatility of all innovations.

The top panel of Table 4 presents the variance-covariance matrix and the standard deviation/correlation matrix of the news terms, estimated as described above. Consistent with previous research, we find that discount-rate news is twice as volatile as cash-flow news.

The interesting new results in this table concern the variance news term N_V. First, it is about as volatile as the discount-rate news. Second, it is highly negatively correlated with cash-flow news: as one might expect from the literature on the "leverage effect" (Black 1976, Christie 1982), news about low cash flows is associated with news about higher future volatility. Third, N_V correlates positively with discount-rate news, indicating that news of high volatility tends to coincide with news of high future real returns. This correlation has been called the "volatility feedback effect" (Campbell and Hentschel 1992, Calvet and Fisher 2007). Both these correlations contribute to a strong negative correlation between volatility shocks and contemporaneous market returns (French, Schwert, and Stambaugh 1987).

The lower right panel of Table 4 reports the decomposition of the vector of innovations $\sigma_t^2 u_{t+1}$ into the three terms $N_{CF,t+1}, N_{DR,t+1}$ and $N_{V,t+1}$. As shocks to $EVAR$ are just a linear combination of shocks to the underlying state variables, which includes $RVAR$, we "unpack" $EVAR$ to express the news terms as a function of r_M, PE, TY, VS, DEF, and $RVAR$. The panel shows that innovations to $RVAR$ are mapped almost one-to-one to news about future volatility. However, several of the other state variables also drive news about volatility. We find that innovations in DEF and VS are associated with news of higher future volatility. Finally, a positive shock to PE (with no change in returns) corresponds to a negative shock to earnings and predicts higher future volatility.

Figure 4 plots the smoothed series for N_{CF}, $-N_{DR}$ and N_V using an exponentially-weighted moving average with a quarterly decay parameter of 0.08. This decay parameter implies a half-life of six years. The pattern of N_{CF} and $-N_{DR}$ we find is consistent with previous research. As a consequence, we focus on the smoothed series for market variance news. There is considerable time variation in N_V, and in particular we find episodes of news of high future volatility during the Great Depression and just before the beginning of World War 2, followed by a period of little news until the late 1960s. From then on, periods of positive volatility news alternate with periods of negative volatility news in cycles of 3 to 5 years. Spikes in news about future volatility are found in the early 1970s (following the oil shocks), in the late 1970s and again following the 1987 crash of the stock market. The late 1990s are characterized by strongly negative news about future returns, and at the same time higher expected future volatility. The recession of the late 2000s is instead characterized by a strong negative cash-flow news, together with a spike in volatility of the highest magnitude in our sample. The recovery from the financial crisis has brought positive cash-flow news together with news about lower future volatility.

3.3 Test assets

In addition to the six VAR state variables, our analysis also requires returns on a cross section of test assets. We construct three sets of portfolios to use as test assets. Our primary cross section consists of the excess returns on the 25 ME- and BE/ME-sorted portfolios, studied in Fama and French (1993), extended in Davis, Fama, and French (2000), and made available by Professor Kenneth French on his web site.[9]

[9] http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/

Daniel and Titman (1997, 2012) and Lewellen, Nagel, and Shanken (2010) point out that it can be misleading to test asset pricing models using only portfolios sorted by characteristics known to be related to average returns, such as size and value. In particular, characteristics-sorted portfolios are likely to show some spread in betas identified as risk by almost any asset pricing model, at least in sample. When the model is estimated, a high premium per unit of beta will fit the large variation in average returns. Thus, at least when premia are not constrained by theory, an asset pricing model may spuriously explain the average returns to characteristics-sorted portfolios.

To alleviate this concern, we follow the advice of Daniel and Titman (1997, 2012) and Lewellen, Nagel, and Shanken (2010) and construct a second set of six portfolios double-sorted on past risk loadings to market and variance risk. First, we run a loading-estimation regression for each stock in the CRSP database where $r_{i,t}$ is the log stock return on stock i for month t.

$$\sum_{j=1}^{3} r_{i,t+j} = b_0 + b_{r_M} \sum_{j=1}^{3} r_{M,t+j} + b_{\Delta VAR} \sum_{j=1}^{3} \Delta VAR_{t+j} + \varepsilon_{i,t+3}$$

We calculate ΔVAR as a weighted sum of changes in the VAR state variables. The weight on each change is the corresponding value in the linear combination of VAR shocks that defines news about market variance. We choose to work with changes rather than shocks as this allows us to generate pre-formation loading estimates at a frequency that is different from our VAR. Namely, though we estimate our VAR using calendar-quarter-end data, our approach allows a stock's loading estimates to be updated at each interim month.

The regression is reestimated from a rolling 36-month window of overlapping observations for each stock at the end of each month. Since these regressions are estimated from stock-level instead of portfolio-level data, we use quarterly data to minimize the impact of infrequent trading. With loading estimates in hand, each month we perform a two-dimensional sequential sort on market beta and EVAR beta. First, we form three groups by sorting stocks on $\widehat{b}_{r_M}$. Then, we further sort stocks in each group to three portfolios on $\widehat{b}_{\Delta VAR}$ and record returns on these nine value-weight portfolios. The final set of risk-sorted portfolios are the two sets of three $\widehat{b}_{r_M}$ portfolios within the extreme $\widehat{b}_{\Delta VAR}$ groups. To ensure that the average returns on these portfolio strategies are not influenced by various market-microstructure issues plaguing the smallest stocks, we exclude the five percent of stocks with the lowest ME from each cross-section and lag the estimated risk loadings by a month in our sorts.

In the empirical analysis, we consider two main subsamples: early (1936:3-1963:3) and modern (1963:4-2010:4) due to the findings in Campbell and Vuolteenaho (2004) of dramatic differences in the risks of these portfolios between the early and modern period. The first subsample is shorter than that in Campbell and Vuolteenaho (2004) as we require each of the 25 portfolios to have at least two stocks as of the time of formation in June.

Finally, we generate a parsimonious cross section of option, bond, and equity returns for the 1986:1-2010:4 time period based on the findings in Fama and French (1993) and Coval and Shumway (2001). In particular, we use the S&P 100 index straddle returns studied by Coval and Shumway.[10] We also include a proxy for the risky bond factor ($DEFRET$) of Fama and French (1993) which we define as the difference in return between the Barclays Capital High Yield Bond Index and the Barclays Capital Investment Grade Bond Index. When pricing the straddle and risky bond return series, we include the returns on the market ($RMRF$), size (SMB), and value (HML) equity factors of Fama and French (1993) as they argue these factors do a good job describing the cross section of average equity returns.

4 Measuring and Pricing Cash-flow, Discount-Rate, and Risk Betas

4.1 Beta Measurement

We next examine the validity of an unconditional version of the first-order condition in equation (14). We modify equation (14) in three ways. First, we use simple expected returns on the left-hand side to make our results easier to compare with previous empirical studies. Second, we condition down equation (14) to avoid having to estimate all required conditional moments. Finally, we cosmetically multiply and divide all three covariances by the sample variance of the unexpected log real return on the market portfolio. By doing so, we can express our pricing equation in terms of betas, facilitating comparison to previous research. These modifications result in the following asset-pricing equation

$$E[R_i - R_f] = \gamma\sigma_M^2\beta_{i,CF_M} + \sigma_M^2\beta_{i,DR_M} - \frac{1}{2}\omega\sigma_M^2\beta_{i,V_M}, \tag{25}$$

where

$$\begin{aligned}
\beta_{i,CF_M} &\equiv \frac{Cov(r_{i,t}, N_{CF,t})}{Var(r_{M,t} - E_{t-1}r_{M,t})}, \\
\beta_{i,DR_M} &\equiv \frac{Cov(r_{i,t}, -N_{DR,t})}{Var(r_{M,t} - E_{t-1}r_{M,t})}, \\
\text{and } \beta_{i,V_M} &\equiv \frac{Cov(r_{i,t}, N_{V,t})}{Var(r_{M,t} - E_{t-1}r_{M,t})}.
\end{aligned}$$

We price the average excess returns on our test assets using the unconditional first-order condition in equation (25) and the quadratic relationship between the parameters ω and γ

[10] Specifically, the series we study includes only those straddle positions where the difference between the options' strike price and the underlying price is between 0 and 5. We thank Josh Coval and Tyler Shumway for providing their updated data series to us.

given by (24). As a first step, we estimate cash-flow, discount-rate, and variance betas using the fitted values of the market's cash flow, discount-rate, and variance news estimated in the previous section. Specifically, we estimate simple WLS regressions of each portfolio's log returns on each news term, weighting each time-$t+1$ observation pair by $(EVAR_t)^{-1}$. We then scale the regression loadings by the ratio of the sample variance of the news term in question to the sample variance of the unexpected log real return on the market portfolio to generate estimates for our three-beta model.

Characteristic-sorted test assets

Table 5 Panel A shows the estimated betas for the 25 size- and book-to-market portfolios over the 1936-1963 period. The portfolios are organized in a square matrix with growth stocks at the left, value stocks at the right, small stocks at the top, and large stocks at the bottom. At the right edge of the matrix we report the differences between the extreme growth and extreme value portfolios in each size group; along the bottom of the matrix we report the differences between the extreme small and extreme large portfolios in each BE/ME category. The top matrix displays post-formation cash-flow betas, the middle matrix displays post-formation discount-rate betas, while the bottom matrix displays post-formation variance betas. In square brackets after each beta estimate we report a standard error, calculated conditional on the realizations of the news series from the aggregate VAR model.

In the pre-1963 sample period, value stocks have both higher cash-flow and higher discount-rate betas than growth stocks. An equal-weighted average of the extreme value stocks across size quintiles has a cash-flow beta 0.13 higher than an equal-weighted average of the extreme growth stocks. The difference in estimated discount-rate betas is also 0.13 and in the same direction. Similar to value stocks, small stocks have higher cash-flow betas and discount-rate betas than large stocks in this sample (by 0.13 and 0.25, respectively, for an equal-weighted average of the smallest stocks across value quintiles relative to an equal-weighted average of the largest stocks). These differences are extremely similar to those in Campbell and Vuolteenaho (2004), despite the exclusion of the 1929-1936 subperiod, the replacement of the excess log market return with the log real return, and the use of a richer, heteroskedastic VAR.

The new finding in Table 5 Panel A is that value stocks and small stocks are also riskier in terms of volatility betas. An equal-weighted average of the extreme value stocks across size quintiles has a volatility beta 0.39 lower than an equal-weighted average of the extreme growth stocks. Similarly, an equal-weighted average of the smallest stocks across value quintiles has a volatility beta that is 0.36 lower than an equal-weighted average of the largest stocks. In summary, value and small stocks were unambiguously riskier than growth and large stocks over the 1936-1963 period.

Table 6 Panel A reports the corresponding estimates for the post-1963 period. As documented in this subsample by Campbell and Vuolteenaho (2004), value stocks still have slightly higher cash-flow betas than growth stocks, but much lower discount-rate betas. Our new finding here is that value stocks continue to have much lower volatility betas, and the spread in volatility betas is even greater than in the early period. The volatility beta for the equal-weighted average of the extreme value stocks across size quintiles is 0.52 lower than the volatility beta of an equal-weighted average of the extreme growth stocks, a difference that is more than 30% higher than the corresponding difference in the early period.

These results imply that in the post-1963 period where the CAPM has difficulty explaining the low returns on growth stocks relative to value stocks, growth stocks hedge two key aspects of the investment opportunity set. Consistent with Campbell and Vuolteenaho (2004), growth stocks hedge news about future real stock returns. The novel finding of this paper is that growth stocks also hedge news about the variance of the market return.

Risk-sorted test assets

Table 5 Panel B shows the estimated betas for the six risk-sorted portfolios over the 1936-1963 period. The portfolios are organized in a rectangular matrix with low CAPM beta stocks at the left, high CAPM beta stocks at the right, low volatility beta stocks at the top, and high volatility beta stocks at the bottom. At the right edge of the matrix we report the differences between the high CAPM beta and the low CAPM beta portfolios in each volatility beta group; along the bottom of the matrix we report the differences between the high volatility beta and the low volatility beta portfolios in each CAPM beta category. As in Panel A, the top matrix displays post-formation cash-flow betas, the middle matrix displays post-formation discount-rate betas, while the bottom matrix displays post-formation volatility betas.

In the pre-1963 sample period, high CAPM beta stocks have both higher cash-flow and higher discount-rate betas than low CAPM beta stocks. An equal-weighted average of the high CAPM beta stocks across the two volatility beta categories has a cash-flow beta 0.19 higher than an equal-weighted average of the low CAPM beta stocks. The difference in estimated discount-rate betas is 0.31 and in the same direction. Similar to high CAPM beta stocks, low volatility beta stocks have higher cash-flow betas and discount-rate betas than high volatility beta stocks in this subsample (by 0.06 and 0.13, respectively, for an equal-weighted average of the low volatility beta stocks across the three CAPM beta categories relative to a corresponding equal-weighted average of the high volatility beta stocks).

High CAPM beta stocks and low volatility beta stocks are also riskier in terms of volatility betas. An equal-weighted average of the high CAPM beta stocks across volatility beta categories has a post-formation volatility beta 0.37 lower than an equal-weighted average of the low CAPM beta stocks. Similarly, an equal-weighted average of the low volatility beta stocks across CAPM beta categories has a post-formation volatility beta that is 0.16 lower than an equal-weighted average of the high volatility beta stocks. In summary, high CAPM beta and low volatility beta stocks were unambiguously riskier than low CAPM beta and high volatility beta stocks over the 1936-1963 period.

Table 6 Panel B shows the estimated betas for the six risk-sorted portfolios over the post-1963 period. In the modern period, high CAPM beta stocks again have higher cash-flow and higher discount-rate betas than low CAPM beta stocks. An equal-weighted average of the high CAPM beta stocks across the two volatility beta categories has a cash-flow beta 0.07 higher than an equal-weighted average of the low CAPM beta stocks. The difference in estimated discount-rate betas is 0.56 and in the same direction. However, high CAPM beta stocks are no longer riskier in terms of volatility betas. Now, an equal-weighted average of the high CAPM beta stocks across the two volatility beta categories has a post-formation variance beta 0.26 higher than a corresponding equal-weighted average of the low CAPM beta stocks. Since, in the three-beta model, covariation with aggregate volatility has a negative premium, the three-beta model can potentially explain why stocks with high past CAPM betas have offered relatively little extra return, at least in the modern period.

In the post-1963 period, sorts on volatility beta continue to generate economically and statistically significant spread in post-formation volatility beta. An equal-weighted average of low volatility beta stocks across the three CAPM beta categories has a post-formation volatility beta that is 0.16 lower than the post-formation volatility beta of a corresponding equal-weighted average of high volatility beta stocks. However, sorts on volatility beta generate very little spread in cash-flow betas or discount-rate betas in the post-1963 period.

Non-equity test assets

Finally, Table 6 Panel C reports the three ICAPM betas of the S&P 100 index straddle position analyzed in Coval and Shumway (2001) along with the corresponding ICAPM betas of the three equity factors and the default bond factor of Fama and French (1993) over the period 1986:1 - 2010:4. Consistent with the nature of a straddle bet, we find that the straddle has a very large volatility beta of 1.27 along with a large negative discount-rate beta of -1.57 and a large (relatively speaking) negative cash-flow beta of -0.14. As one would expect, the betas of the Fama-French equity factors are consistent with the findings for the size- and book-to-market-sorted portfolios in Table 6 Panel B. Finally, the risky bond factor $DEFRET$ has a cash-flow beta of 0.06, a discount-rate beta of 0.20, and a volatility beta of -0.31. The sign of $DEFRET$'s volatility beta is consistent with the fact that risky corporate debt is short the option to default.

4.2 Beta Pricing

We next turn to pricing the cross section with these three ICAPM betas. We evaluate the performance of five asset-pricing models: 1) the traditional CAPM that restricts cash-flow and discount-rate betas to have the same price of risk and sets the price of variance risk equal to zero; 2) the two-beta intertemporal asset pricing model of Campbell and Vuolteenaho (2004) that restricts the price of discount-rate risk to equal the variance of the market return, 3) our three-beta intertemporal asset pricing model that restricts the price of discount-rate risk to equal the variance of the market return and constrains the price of cash-flow and variance risk to be related by equation (24), with $\rho = 0.95$ per year; 4) a partially-constrained three-beta model that restricts the price of discount-rate risk to equal the variance of the market return but freely estimates the other two risk prices (effectively decoupling γ and ω), and 5) an unrestricted three-beta model that allows free risk prices for cash-flow, discount-rate, and volatility betas. Each model is estimated in two different forms: one with a restricted zero-beta rate equal to the Treasury-bill rate, and one with an unrestricted zero-beta rate following Black (1972).

Characteristic-sorted test assets

Table 7 reports results for the early sample period 1936-1963, using 25 size- and book-to-market-sorted portfolios as test assets. The table has ten columns, two specifications for each of our five asset pricing models. The first 16 rows of Table 7 are divided into four sets of four rows. The first set of four rows corresponds to the zero-beta rate (in excess of the Treasury-bill rate), the second set to the premium on cash-flow beta, the third set to the premium on discount-rate beta, and the fourth set to the premium on volatility beta. Within each set, the first row reports the point estimate in fractions per quarter, and the second row annualizes this estimate, multiplying by 400 to aid in interpretation. These parameters are estimated from a cross-sectional regression

$$\overline{R}_i^e = g_0 + g_1\widehat{\beta}_{i,CF_M} + g_2\widehat{\beta}_{i,DR_M} + g_3\widehat{\beta}_{i,V_M} + e_i, \tag{26}$$

where a bar denotes time-series mean and $\overline{R}_i^e \equiv \overline{R}_i - \overline{R}_{rf}$ denotes the sample average simple excess return on asset i. The third and fourth rows present two alternative standard errors of the monthly estimate, described below.

Below the premia estimates, we report the R^2 statistic for a cross-sectional regression of average returns on our test assets onto the fitted values from the model. We also report a composite pricing error, computed as a quadratic form of the pricing errors. The weighting matrix in the quadratic form is a diagonal matrix with the inverse of the sample test asset return volatilities on the main diagonal.

Standard errors are produced with a bootstrap from 2,500 simulated realizations. Our bootstrap experiment samples test-asset returns and first-stage VAR errors, and uses the first-stage and second-stage WLS VAR estimates in Table 2 to generate the state-variable data.[11] We partition the VAR errors and test-asset returns into two groups, one for 1936 to 1963 and another for 1963 to 2010, which enables us to use the same simulated realizations in subperiod analyses. The first set of standard errors (labeled A) conditions on estimated news terms and generates betas and return premia separately for each simulated realization, while the second set (labeled B) also estimates the first-stage and second-stage VAR and the news terms separately for each simulated realization. Standard errors B thus incorporate

[11] When simulating the bootstrap, we drop realizations which would result in negative $RVAR$ and redraw. On a related note, our constrained WLS procedure forces our forecast of $RVAR$ and $EVAR$ to be within a positive range.

the considerable additional sampling uncertainty due to the fact that the news terms as well as betas are generated regressors.

Two alternative 5-percent critical values for the composite pricing error are produced with a bootstrap method similar to the one we have described above, except that the test-asset returns are adjusted to be consistent with the pricing model before the random samples are generated. Critical values A condition on estimated news terms, while critical values B take account of the fact that news terms must be estimated.

Finally, Table 7 reports the implied risk-aversion coefficient, γ, which can be recovered as g_1/g_2, as well as the sensitivity of news about risk to news about market variance, ω, which can be recovered as $-2 * g_3/g_2$. The three-beta ICAPM estimates are constrained so that both γ and the implied ω are strictly positive.

Table 7 shows that in the 1936-1963 period, the restricted three-beta model explains the cross-section of stock returns reasonably well. The cross-sectional R^2 statistics are almost 60% for both forms of this model. Both the Sharpe-Lintner and Black versions of the CAPM do a slightly poorer job describing the cross section (R^2 statistics are 52% and 53% respectively). The two-beta ICAPM of Campbell and Vuolteenaho (2004) performs slightly better than the CAPM and slightly worse than the volatility ICAPM. None of the theoretically-motivated models considered are rejected by the data based on the composite pricing test. Consistent with the claim that the three-beta model does a good job describing the cross-section, Table 7 shows that the constrained and the unrestricted factor model barely improve pricing relative to the three-beta ICAPM.

Figure 5 provides a visual summary of these results. The figure plots the predicted average excess return on the horizontal axis and the actual sample average excess return on the vertical axis. In summary, we find that the three-beta ICAPM improves pricing relative to both the Sharpe-Lintner and Black versions of the CAPM.

This success is due in part to the inclusion of volatility betas in the specification. For the Black version of the three-beta ICAPM, the spread in volatility betas across the 25 size- and book-to-market-sorted portfolios generates an annualized spread in average returns of 2.77% compared to a comparable spread of 7.61% and 2.45% for cash-flow and discount-rate betas. Variation in volatility betas accounts for 5% of the variation in explained returns compared to 36% and 4% for cash-flow and discount-rate betas respectively. The remaining 55% of the explained variation in average returns is due of course to the covariation among the three types of betas.

Results are very different in the 1963-2010 period. Table 8 shows that in this period, both versions of the CAPM do a very poor job of explaining cross-sectional variation in average returns on portfolios sorted by size and book-to-market. When the zero-beta rate is left as a free parameter, the cross-sectional regression picks a negative premium for the CAPM beta and implies an R^2 of slightly under 6%. When the zero-beta rate is constrained to the risk-free rate, the CAPM R^2 falls to roughly -39%. Both versions of the static CAPM are easily rejected at the five-percent level by both sets of critical values.

In the modern period, the unconstrained zero-beta rate version of the two-beta Campbell and Vuolteenaho (2004) model does a better job describing the cross section of average returns than the CAPM. However, the implied coefficient of risk aversion, 19.90, is arguably extreme.

The three-beta model with the restricted zero-beta rate also does a poor job explaining cross-sectional variation in average returns across our test assets. However, if we continue to restrict the risk price for discount-rate and variance news but allow an unrestricted zero-beta rate, the explained variation increases to roughly 72%, more than twice the R^2 of the corresponding two-beta ICAPM. The estimated risk price for cash-flow beta is an economically reasonable 28 percent per year with an implied coefficient of relative risk aversion of 9.13. Both versions of our intertemporal CAPM with stochastic volatility are not rejected at the 5-percent level by either set of critical values.

Figure 6 provides a visual summary of these results. For the Black version of the three-beta ICAPM, spread in volatility betas across the 25 size- and book-to-market-sorted portfolios generates an annualized spread in average returns of 7.84% compared to a comparable spread of 3.48% and 2.20% for cash-flow and discount-rate betas. Variation in volatility betas accounts for 99% of the variation in explained returns compared to 14% for cash-flow betas as well as 6% for discount-rate betas. Covariation among the three types of betas is responsible for the remaining -19% of explained variation in average returns.

The relatively poor performance of the risk-free rate version of the three-beta ICAPM is due to the derived link between γ and ω. To show this, Figure 7 provides two contour plots (one each for the risk-free and zero-beta rate versions of the model in the top and bottom panels of the figure respectively) of the R^2 resulting from combinations of (γ,ω) ranging from (0,0) to (40,16). On the same figure we also plot the relation between γ and ω derived in equation (24). The top panel of Figure 7 shows that even with the intercept restricted to zero, R^2's are as high as 70% for some combinations of (γ,ω). Unfortunately, as the plot shows, these combinations do not coincide with the curve implied by equation (24). Once the zero-beta rate is unconstrained, the contours for R^2's greater than 70% cover a much larger area of the plot and coincide nicely with the ICAPM relation of equation (24).

Consistent with the contour plots of Figure 7, the pricing results in Table 8 based on the partially-constrained factor model further confirms that the link between γ and ω is responsible for the poor fit of the restricted zero-beta rate version of the three-beta ICAPM in the modern period. When removing the constraint linking γ and ω but leaving the constraint on the discount-rate beta premium in place, the R^2 increases from -118% to 76%. Nevertheless, the risk prices for γ and ω remain economically large and of the right sign.

Risk-sorted test assets

We confirm that the success of the three-beta ICAPM is robust by expanding the set of test portfolios beyond the 25 size- and book-to-market-sorted portfolios. First, we show that our three-beta model not only describes the cross section of characteristics-sorted portfolios but also can explain the average returns on risk-sorted portfolios. We examine risk-sorted portfolios as Daniel and Titman (1997, 2012) and Lewellen, Nagel, and Shanken (2010) argue that asset-pricing tests using only portfolios sorted by characteristics known to be related to average returns, such as size and value, can be misleading due to the low-dimensional factor structure of the 25 size and book-to-market-sorted portfolios.

Table 9 prices the six risk-sorted portfolios described in Table 5 Panel B in conjunction with six of the 25 size- and book-to-market-sorted portfolios of Table 5 Panel A (the low, medium, and high BE/ME portfolios within the small and large ME quintiles). We continue to find that the three-beta ICAPM improves pricing relative to both the Sharpe-Lintner and Black versions of the CAPM. Moreover, the relatively high R^2 (61%) is not disproportionately due to characteristics-sorted portfolios as the R^2 for the risk-sorted subset (57%) is comparable to the R^2 for the characteristics-sorted subset (58%). Figure 8 shows this success graphically.

Table 10 prices the cross section of characteristics- and risk-sorted portfolios in the modern period. We find that the zero-beta rate three-beta ICAPM is not rejected by the data while both versions of the CAPM are rejected. Again, the relatively high R^2 for the zero-beta rate version of the volatility ICAPM (75%) is not disproportionately due to characteristics-sorted portfolios as the R^2 for the risk-sorted subset (69%) is comparable to the R^2 for the characteristics-sorted subset (84%). Figure 9 provides a graphically summary of these results.

Non-equity test assets

We also show that our three-beta model can help explain average returns on non-equity portfolios designed to be highly correlated with aggregate volatility risk, namely the S&P 100 index straddles of Coval and Shumway (2001). We first calculate the expected return on straddle portfolio based on the estimates of the zero-beta rate volatility ICAPM in Table 8. The contributions to expected quarterly return from the straddle's cash-flow, discount-rate, and volatility betas are -0.96%, -1.19%, and -3.73% respectively. As the average quarterly realized return on the straddle is -23.61%, an equity-based estimate of the three-beta model explains roughly 25% of the realized straddle premium.

Table 11 shows that our intertemporal CAPM with stochastic volatility is not rejected at the 5-percent level when we price the joint cross-section of equity, bond, and straddle returns. The implied risk aversion coefficient (roughly 14 for both the risk-free and zero-beta rate implementations of the model) is high but not unreasonable. In sharp contrast, the CAPM is strongly rejected. Though the two-beta ICAPM is not rejected, the required risk aversion is too extreme (over 78 for both versions of the model) to be realistic.

Summary of US financial history

Figure 10 (third panel) plots the time-series of the smoothed combined shock $\gamma N_{CF} - N_{DR} - \frac{1}{2}\omega N_V$ based on the estimate of the zero-beta model for the modern period (Table 7). The correlation of this shock with the associated N_{CF} is 0.90. Similarly, the correlation of this shock with the associated N_{DR} is 0.26. Finally, the correlation of this shock with the associated N_V is -0.76. Figure 10 also plots the corresponding smoothed shock series for the CAPM ($N_{CF} - N_{DR}$) and for the two-beta ICAPM ($\gamma N_{CF} - N_{DR}$). The two-beta model shifts the history of good and bad times relative to the CAPM, as emphasized by Campbell, Giglio, and Polk (2011). The model with stochastic volatility further accentuates that periods with high market volatility, such as the 1930s and the late 2000s, are particularly hard times for long-term investors.

Finally, we note that our key findings are robust to a variety of methodological changes. These changes include forecasting excess rather than real returns in the VAR, using OLS rather than WLS, using a unconstrained WLS approach rather than a constrained WLS approach, and setting the parameter ρ to any value between 0.88 and 0.97.

5 Conclusion

We extend the approximate closed-form intertemporal capital asset pricing model of Campbell (1993) to allow for stochastic volatility. Our model recognizes that an investor's investment opportunities may deteriorate either because expected stock returns decline or because the volatility of stock returns increases. A conservative long-term investor will wish to hedge against both types of changes in investment opportunities; thus, a stock's risk is determined not only by its beta with unexpected market returns and news about future returns (or equivalently, news about market cash flows and discount rates), but also by its beta with news about future market volatility. Although our model has three dimensions of risk, the prices of all these risks are determined by a single free parameter, the coefficient of relative risk aversion.

Our implementation models the return on the aggregate stock market as one element of a vector autoregressive (VAR) system; the volatility of all shocks to the VAR is another element of the system. We show that the negative post-1963 CAPM alphas of growth stocks are justified because these stocks hedge long-term investors against both declining expected stock returns, and increasing volatility. The addition of volatility risk to the model helps it to deliver a moderate, economically reasonable value of risk aversion.

Our empirical work is limited in one important respect. We test only the unconditional implications of the model and do not evaluate its conditional implications. A full conditional test is likely to be a challenging hurdle for the model. To see why, recall that we assume a rational long-term investor always holds 100% of his or her assets in equities. However, time-variation in real stock returns generally gives the long-term investor an incentive to shift the relative weights on cash and equity, unless real interest rates and market volatility move in exactly the right way to make the equity premium proportional to market volatility. Although we do not explicitly test whether this is the case, previous work by Campbell (1987) and Harvey (1989, 1991) rejects this proportionality restriction.

One way to support the assumption of constant 100% equity investment is to invoke binding leverage constraints. Indeed, in the modern sample, the Black (1972) version of our three-beta model is consistent with this interpretation as the estimated difference between the zero-beta and risk-free rates is positive, statistically significant, and economically large. However, the risk aversion coefficient we estimate may be too large to explain why leverage constraints should bind.

Nevertheless, our model does directly answer the interesting microeconomic question: Are there reasonable preference parameters that would make a long-term investor, constrained to invest 100% in equity, content to hold the market rather than tilting towards value stocks or other high-return stock portfolios? Our answer is clearly yes.

Appendix

Deriving the equation for ω

Here we show how to solve for the unknown parameter ω as discussed in section 2. We start from the definition of ω

$$\begin{aligned}
\omega\sigma_t^2 &= \mathrm{Var}_t\left[m_{t+1}+r_{t+1}\right] \\
&= \mathrm{Var}_t\left[\frac{\theta}{\psi}h_{t+1}+(1-\gamma)r_{t+1}\right] \\
&= \mathrm{Var}_t\left[\frac{\theta}{\psi}\left((\psi-1)N_{DR,t+1}+\frac{1}{2}\frac{\psi}{\theta}\omega N_{V,t+1}\right)+(1-\gamma)r_{t+1}\right] \\
&= \mathrm{Var}_t\left[(1-\gamma)N_{DR,t+1}+\frac{1}{2}\omega N_{V,t+1}+(1-\gamma)r_{t+1}\right] \\
&= \mathrm{Var}_t\left[(1-\gamma)N_{CF,t+1}+\frac{1}{2}\omega N_{V,t+1}\right] \\
&= (1-\gamma)^2\mathrm{Var}_t\left[N_{CF_{t+1}}\right]+\omega(1-\gamma)\mathrm{Cov}_t\left[N_{CF_{t+1}},N_{V_{t+1},}\right]+\frac{\omega^2}{4}\mathrm{Var}_t\left[N_{V_{t+1}}\right],
\end{aligned}$$

deriving equation (19). Since cash flow and volatility news can be expressed in terms of the VAR parameters as

$$\begin{aligned}
N_{V,t+1} &= e_2'(I-\rho\Gamma)^{-1}\sigma_t u_{t+1} \\
N_{CF,t+1} &= (e_1'+e_1'\rho\Gamma(I-\rho\Gamma)^{-1})\sigma_t u_{t+1}
\end{aligned}$$

we can define the covariance matrix of VAR shocks as $\mathbf{\Sigma}$ =$\mathrm{Var}_t\left[u_{t+1}\right]$ =$\mathrm{Var}[u_{t+1}]$ and the error-to-news vectors x_{CF} and x_V, defined in equations (20) and (21), to write ω as the solution to

$$0=\omega^2\frac{1}{4}x_V\mathbf{\Sigma}x_V'-\omega\left(1-(1-\gamma)\,x_{CF}\mathbf{\Sigma}x_V'\right)+(1-\gamma)^2\,x_{CF}\mathbf{\Sigma}x_{CF}'$$

as was presented in equation (22).

Selecting the correct root of the quadratic equation

The equation defining ω will generally have two solutions

$$\omega=\frac{1-(1-\gamma)\,x_{CF}\mathbf{\Sigma}x_V'\pm\sqrt{\left(1-(1-\gamma)\,x_{CF}\mathbf{\Sigma}x_V'\right)^2-(1-\gamma)^2\left(x_V\mathbf{\Sigma}x_V'\right)\left(x_{CF}\mathbf{\Sigma}x_{CF}'\right)}}{\frac{1}{2}x_V\mathbf{\Sigma}x_V'}.$$

As was discussed in the paper, this is an artifact of the loglinear approximation. While the (approximate) Euler equation holds for both roots, the correct solution is the one with the negative sign on the radical shown in equation (23).

This can be confirmed from numerical computation, and it can also be easily seen by observing the behavior of the solutions in the limit as volatility news goes to zero and the model become homoskedastic. With the false solution, ω becomes infinitely large as $x_V \to 0$. This corresponds to the log value of invested wealth going to negative infinity. On the other hand, we can use the correct solution for ω converges to $(1-\gamma)^2 x_{CF}\mathbf{\Sigma}x'_{CF}$. This is what we would expect, since in that case $\omega = \frac{1}{\sigma_t}\text{Var}_t\left[(1-\gamma)N_{CF,t+1}\right]$.

An approximation for ω

As discussed in section 2, we will not find a real solution for ω if

$$(1-\gamma)^2 (x_V\mathbf{\Sigma}x'_V)(x_{CF}\mathbf{\Sigma}x'_{CF}) > (1-(1-\gamma)x_{CF}\mathbf{\Sigma}x'_V)^2 .$$

This an unfortunate artifact of the loglinearization approach, as the conditional variance defining ω =Var$_t\left[\frac{m_{t+1}+r_{t+1}}{\sigma_t}\right]$ will be real and finite for the true stochastic discount factor, m_{t+1}. We propose an alternative approach that will allow us to approximate ω even when $\gamma > 6.4$, the region where there is no real solution given our estimated VAR parameters.

We start from the definition of ω

$$\begin{aligned}\omega &= \frac{1}{\sigma_t^2}\text{Var}_t\left[(1-\gamma)N_{CF,t+1} + \frac{1}{2}\omega N_{V,t+1}\right] \\ &= \omega^2\frac{1}{4}x_V\mathbf{\Sigma}x'_V + \omega(1-\gamma)x_{CF}\mathbf{\Sigma}x'_V + (1-\gamma)^2 x_{CF}\mathbf{\Sigma}x'_{CF}\end{aligned}$$

and approximate the variance term on the right hand side so that it is a linear function of ω rather than quadratic. Taking a Taylor approximation about $\omega = 0$

$$\omega \approx (1-\gamma)^2 x_{CF}\mathbf{\Sigma}x'_{CF} + (1-\gamma)x_{CF}\mathbf{\Sigma}x'_V\omega$$

where the first term on the right hand side is the traditional value that we would see in the homoskedastic case, as pointed out in the text. The second term is the additional effect coming from stochastic volatility. Now, solving for ω, we generate the approximation used in the empirical analysis.

$$\omega \approx \frac{(1-\gamma)^2 x_{CF}\mathbf{\Sigma}x'_{CF}}{1-(1-\gamma)x_{CF}\mathbf{\Sigma}x'_V}$$

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* 63:2997–3030.

Andersen, Torben G., Tim Bollerslev, Francis X. Diebold, and Paul Labys, 2003, "Modeling and Forecasting Realized Volatility", *Econometrica* 71:579–625.

Ang, Andrew, Robert J. Hodrick, Yuhang Xing, and Xiaoyan Zhang, 2006, "The Cross-Section of Volatility and Expected Returns", *Journal of Finance* 61:259–299.

Bansal, Ravi and Amir Yaron, 2004, "Risks for the Long Run", *Journal of Finance* 59:1481–1509.

Bansal, Ravi, Dana Kiku and Amir Yaron, 2012, "An Empirical Evaluation of the Long-Run Risks Model for Asset Prices", *Critical Finance Review* 1:183–221.

Bansal, Ravi, Dana Kiku, Ivan Shaliastovich and Amir Yaron, 2011, "Volatility, the Macroeconomy and Asset Prices", upublished paper, Duke University and University of Pennsylvania.

Barndorff-Nielsen, Ole E. and Neil Shephard, 2002, "Econometric Analysis of Realized Volatility and Its Use in Estimating Stochastic Volatility Models", *Journal of the Royal Statistical Society B*, 64(2):253–280.

Beeler, Jason and John Y. Campbell, 2012, "The Long-Run Risks Model and Aggregate Asset Prices: An Empirical Assessment", *Critical Finance Review* 1:141-182.

Black, Fisher, 1972, "Capital Market Equilibrium with Restricted Borrowing", *Journal of Business* 45 444-454.

Black, Fisher, 1976, "Studies of Stock Price Volatility Changes", *Proceedings of the 1976 Meetings of the American Statistical Association, Business and Economic Statistics Section,* Washington 177-181.

Breeden, Douglas T., 1979, "An Intertemporal Asset Pricing Model with Stochastic Consumption and Investment Opportunities", *Journal of Financial Economics* 7:265–296.

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.

Campbell, John Y., 1993, "Intertemporal Asset Pricing Without Consumption Data", *American Economic Review* 83:487–512.

Campbell, John Y., 1996, "Understanding Risk and Return", *Journal of Political Economy* 104:298–345.

Campbell, John Y., Stefano Giglio, and Christopher Polk, 2011, "Hard Times", London School of Economics working paper.

Campbell, John Y., Stefano Giglio, Christopher Polk, and Robert Turley, 2012, "Internet Appendix to An Intertemporal CAPM with Stochastic Volatility", available online at http://kuznets.fas.harvard.edu/~campbell/papers.html.

Campbell, John Y. 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., Christopher Polk, and Tuomo Vuolteenaho, 2010, "Growth or Glamour? Fundamentals and Systematic Risk in Stock Returns" *Review of Financial Studies* 23:305–344.

Campbell, John Y. and Robert J. Shiller, 1988a, "The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors", *Review of Financial Studies* 1:195–228.

Campbell, John Y. and Robert J. Shiller, 1988b, "Stock Prices, Earnings, and Expected Dividends", *Journal of Finance* 43:661–676.

Campbell, John Y. and Tuomo Vuolteenaho, 2004, "Bad Beta, Good Beta", *American Economic Review* 94:1249–1275.

Chacko, George and Luis M. Viceira, 2005, "Dynamic Consumption and Portfolio Choice with Stochastic Volatility in Incomplete Markets", *Review of Financial Studies* 18:1369–1402.

Chen, Joseph, 2003, "Intertemporal CAPM and the Cross Section of Stock Returns", unpublished paper, University of California Davis.

Chen, Long and Xinlei Zhao, 2009, "Return Decomposition", *Review of Financial Studies* 22:5213–5249.

Christie, Andrew, "The Stochastic Behavior of Common Stock Variances – Value, Leverage, and Interest Rates Effects", *Journal of Financial Economics* 10:407–432.

Cochrane, John H., 2008, "The Dog That Did Not Bark: A Defense of Return Predictability", *Review of Financial Studies* 21:1533-1575.

Cochrane, John H., 2011, "Discount Rates", NBER Working Paper No. 16972, forthcoming *Journal of Finance*.

Coval, Josh, and Tyler Shumway, 2001, "Expected Option Returns", *Journal of Finance* 56:983–1009.

Daniel, Kent and Sheridan Titman, 1997, "Evidence on the Characteristics of Cross-sectional Variation in Common Stock Returns", *Journal of Finance* 52:1–33.

Daniel, Kent and Sheridan Titman, 2012, "Testing Factor-Model Explanations of Market Anomalies", *Critical Finance Review* 1:103–139.

Davis, James L., Eugene F. Fama, and Kenneth R. French, "Characteristics, Covariances, and Average Returns: 1929 to 1997", *Journal of Finance* 55:389–406.

Duffee, Greg, 2005, "Time-Variation in the Covariance Between Stock Returns and Consumption Growth", *Journal of Finance* 60:1673–172.

Engsted, Tom, Thomas Q. Pedersen, and Carsten Tanggaard, 2010, "Pitfalls in VAR Based Return Decompositions: A Clarification", unpublished paper, University of Aarhus.

Epstein, Lawrence and Stanley Zin, 1989, "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework", *Econometrica* 57:937–69.

Eraker, Bjorn, 2008, "Affine General Equilibrium Models", *Management Science* 54:2068–2080.

Eraker, Bjorn and Ivan Shaliastovich, 2008, "An Equilibrium Guide to Designing Affine Pricing Models", *Mathematical Finance* 18:519–543.

Eraker, Bjorn and Wenyu Wang, 2011, "Dynamic Present Values and the Intertemporal CAPM", University of Wisconsin working paper.

French, Ken, William Schwert, and Robert Stambaugh, 1987, "Expected stock returns and volatility", *Journal of Financial Economics* 19:3–29.

Harvey, Campbell, 1989, "Time-Varying Conditional Covariances in Tests of Asset Pricing Models", *Journal of Financial Economics* 24:289–317.

Harvey, Campbell, 1991, "The World Price of Covariance Risk", *Journal of Finance* 46:111–157.

Kandel, Shmuel and Robert Stambaugh, 1991, "Asset Returns and Intertemporal Preferences", *Journal of Monetary Economics* 27:39–71.

Lewellen, Jonathan, Stefan Nagel, and Jay Shanken, 2010, "A Skeptical Appraisal of Asset Pricing Tests", *Journal of Financial Economics* 96:175–194.

Liu, Jun, 2007, "Portfolio Selection in Stochastic Environments", *Review of Financial Studies* 20:1–39.

Merton, Robert C., 1973, "An Intertemporal Capital Asset Pricing Model", *Econometrica* 41:867–87.

Restoy, Fernando, and Philippe Weil, 1998, "Approximate Equilibrium Asset Prices", NBER Working Paper 6611.

Restoy, Fernando, and Philippe Weil, 2011, "Approximate Equilibrium Asset Prices", *Review of Finance* 15, 1–28.

Sohn, Bumjean, 2010, "Stock Market Volatility and Trading Strategy Based Factors", unpublished paper, Georgetown University.

Tauchen, George and Robert Hussey, 1991, "Quadrature-Based Methods for Obtaining Approximate Solutions to Nonlinear Asset Pricing Models", *Econometrica* 59:317–96.

Table 1: Forecasting Realized Variance

The table reports the WLS parameter estimates of a constrained regression forecasting the realized variance ($RVAR$) of within-quarter daily returns on the CRSP value-weight index. The forecasting variables include lagged values of $RVAR$, the log real return on the CRSP value-weight index (r_M), the log ratio of the S&P 500's price to the S&P 500's ten-year moving average of earnings (PE), the term yield spread (TY) in percentage points, measured as the difference between the log yield on the ten-year US constant-maturity bond and the log yield on the three-month US Treasury Bill, the default yield spread (DEF) in percentage points, measured as the difference between the log yield on Moody's BAA bonds and the log yield on Moody's AAA bonds, and the small-stock value-spread (VS), the difference in the log book-to-market ratios of small value and small growth stocks. The small-value and small-growth portfolios are two of the six elementary portfolios constructed by Davis et al. (2000). The forecasted values from this regression are the state variable EVAR used in the second stage of the estimation. Initial WLS weights of each observation are inversely proportional to $RVAR_t$. These weights are then shrunk towards equal weights so that the maximum ratio of actual weights used is bounded by the corresponding historical ratio for the VIX index. Similarly, the regression estimates are constrained to generate fitted values that fall within the historical range of the VIX. The first seven columns report coefficients on the seven explanatory variables, and the remaining column shows the R^2 and F statistics. Newey-West standard errors estimated with four lags are in square brackets. The sample period for the dependent variables is 1926.2-2010.4, 338 quarterly data points.

Variance Forecast

First stage	Constant	$r_{M,t}$	$RVAR_t$	PE_t	TY_t	DEF_t	VS_t	$R^2\%/F$
$RVAR_{t+1}$	-0.037	-0.017	0.298	0.013	-0.002	0.024	0.001	10.09%
	[0.025]	[0.021]	[0.061]	[0.007]	[0.002]	[0.006]	[0.008]	6.19

Table 2: VAR Parameter Estimates

The table shows the WLS parameter estimates for a first-order VAR model including a constant, the log real market return (r_M), forecasted variance ($EVAR$), price-earnings ratio (PE), term yield spread (TY), default yield spread (DEF), and small-stock value spread (VS). Initial weights of each observation are inversely proportional to $EVAR_t$. These weights are then shrunk towards equal weights so that the maximum ratio of actual weights used is bounded by the corresponding historical ratio for the VIX index. Similarly, the regression estimates are constrained to generate fitted values that fall within the historical range of the VIX. In Panel A of the Table, each set of three rows corresponds to a different dependent variable. The first seven columns report coefficients on the seven explanatory variables, and the remaining column shows the R^2 and F statistics. Newey-West standard errors estimated with four lags are in square brackets and bootstrap standard errors in parentheses. Bootstrap standard errors are computed from 2,500 simulated sample realizations. Panel B of the table reports the correlation matrices of both the unscaled and scaled shocks with shock standard deviations on the diagonal, labeled "corr/std.", as well as the autocorrelation matrices of both the unscaled and scaled shocks, labeled "Autocorr." The sample period for the dependent variables is 1926.2-2010.4, 338 quarterly data points.

Panel A: VAR Estimates

Second stage	Constant	$r_{M,t}$	$EVAR_t$	PE_t	TY_t	DEF_t	VS_t	$R^2\%/F$
$r_{M,t+1}$	0.199	0.124	0.664	-0.054	0.007	-0.029	-0.017	4.59%
	[0.071]	[0.056]	[0.518]	[0.021]	[0.006]	[0.024]	[0.020]	2.65
	(0.129)	(0.082)	(0.927)	(0.039)	(0.009)	(0.028)	(0.047)	
$EVAR_{t+1}$	-0.040	-0.004	0.342	0.012	-0.001	0.018	0.005	25.03%
	[0.011]	[0.009]	[0.080]	[0.003]	[0.001]	[0.004]	[0.003]	18.42
	(0.023)	(0.005)	(0.085)	(0.007)	(0.001)	(0.004)	(0.008)	
PE_{t+1}	0.122	0.190	0.569	0.959	0.007	-0.024	-0.004	99.23%
	[0.067]	[0.053]	[0.494]	[0.020]	[0.005]	[0.023]	[0.019]	7142.61
	(0.122)	(0.079)	(0.878)	(0.037)	(0.008)	(0.027)	(0.044)	
TY_{t+1}	-0.046	-0.161	2.911	-0.002	0.851	0.099	0.044	75.94%
	[0.366]	[0.289]	[2.679]	[0.106]	[0.029]	[0.124]	[0.102]	174.08
	(0.526)	(0.374)	(4.011)	(0.160)	(0.039)	(0.125)	(0.198)	
DEF_{t+1}	0.125	-0.448	2.231	-0.033	-0.003	0.865	0.035	71.46%
	[0.163]	[0.129]	[1.193]	[0.047]	[0.013]	[0.055]	[0.045]	138.10
	(0.285)	(0.200)	(1.822)	(0.080)	(0.020)	(0.064)	(0.102)	
VS_{t+1}	0.122	0.066	0.970	-0.010	-0.005	-0.001	0.930	95.51%
	[0.057]	[0.045]	[0.417]	[0.017]	[0.005]	[0.019]	[0.016]	1174.01
	(0.115)	(0.073)	(0.743)	(0.033)	(0.008)	(0.025)	(0.041)	

Panel B: Correlations and Standard Deviations

corr/std	r_M	$EVAR$	PE	TY	DEF	VS
			unscaled			
r_M	0.107	-0.546	0.908	-0.031	-0.493	-0.046
$EVAR$	-0.546	0.015	-0.641	-0.086	0.698	0.126
PE	0.908	-0.641	0.100	-0.016	-0.599	-0.071
TY	-0.031	-0.086	-0.016	0.565	0.020	-0.021
DEF	-0.493	0.698	-0.599	0.020	0.295	0.328
VS	-0.046	0.126	-0.071	-0.021	0.328	0.087
			scaled			
r_M	0.552	-0.523	0.901	-0.078	-0.289	0.045
$EVAR$	-0.523	0.067	-0.587	-0.077	0.646	0.063
PE	0.901	-0.587	0.499	-0.071	-0.382	0.030
TY	-0.078	-0.077	-0.071	3.156	0.027	-0.018
DEF	-0.289	0.646	-0.382	0.027	1.127	0.238
VS	0.045	0.063	0.030	-0.018	0.238	0.505
Autocorr.	$r_{M,t+1}$	$EVAR_{t+1}$	PE_{t+1}	TY_{t+1}	DEF_{t+1}	VS_{t+1}
			unscaled			
$r_{M,t}$	-0.135	0.044	-0.123	0.079	0.149	0.066
$EVAR_t$	0.110	-0.088	0.117	-0.142	-0.264	-0.126
PE_t	-0.144	0.133	-0.200	0.099	0.271	0.115
TY_t	-0.051	0.078	-0.034	-0.117	0.097	0.059
DEF_t	0.191	-0.130	0.218	-0.170	-0.349	-0.169
VS_t	0.030	-0.059	0.026	-0.077	-0.085	-0.089
			scaled			
$r_{M,t}$	-0.003	-0.043	-0.014	0.018	0.016	-0.014
$EVAR_t$	0.060	-0.002	0.070	-0.062	-0.127	-0.063
PE_t	-0.019	0.033	-0.075	0.022	0.093	0.007
TY_t	-0.042	0.054	-0.037	-0.033	0.073	0.031
DEF_t	0.084	-0.056	0.105	-0.084	-0.208	-0.105
VS_t	0.033	-0.064	0.016	-0.022	-0.073	-0.068

Table 3: VAR Specification Test

The table reports the results of regressions forecasting the squared second-stage residuals from the VAR estimated in Table 2 with $EVAR_t$. Newey-West standard errors estimated with four lags are in square brackets. The sample period for the dependent variables is 1926.2-2010.4, 338 quarterly data points.

Heteroskedastic Shocks

Squared, second-stage, unscaled residual	Constant	$EVAR_t$	R^2%
$r_{M_{t+1}}$	-0.009	0.649	23.78%
	[0.007]	[0.253]	
$EVAR_{t+1}$	-0.000	0.014	5.54%
	[0.000]	[0.003]	
PE_{t+1}	-0.010	0.641	23.68%
	[0.007]	[0.268]	
TY_{t+1}	0.115	6.464	3.73%
	[0.071]	[2.464]	
DEF_{t+1}	-0.213	9.532	32.09%
	[0.093]	[3.595]	
VS_{t+1}	0.003	0.152	6.67%
	[0.001]	[0.039]	

Table 4: Cash-flow, Discount-rate, and Variance News for the Market Portfolio
The table shows the properties of cash-flow news (N_{CF}), discount-rate news (N_{DR}), and volatility news (N_V) implied by the VAR model of Table 2. The upper-left section of the table shows the covariance matrix of the news terms. The upper-right section shows the correlation matrix of the news terms with standard deviations on the diagonal. The lower-left section shows the correlation of shocks to individual state variables with the news terms. The lower-right section shows the functions ($\mathbf{e1}'+\mathbf{e1}'\lambda_{DR}$, $\mathbf{e1}'\lambda_{DR}$, $\mathbf{e2}'\lambda_V$) that map the state-variable shocks to cash-flow, discount-rate, and variance news. We define $\lambda_{DR} \equiv \rho\mathbf{\Gamma}(\mathbf{I}-\rho\mathbf{\Gamma})^{-1}$ and $\lambda_V \equiv \rho(\mathbf{I}-\rho\mathbf{\Gamma})^{-1}$, where $\mathbf{\Gamma}$ is the estimated VAR transition matrix from Table 2 and ρ is set to 0.95 per annum. r_M is the log real return on the CRSP value-weight index. $RVAR$ is the realized variance of daily returns on the CRSP value-weight index. PE is the log ratio of the S&P 500's price to the S&P 500's ten-year moving average of earnings. TY is the term yield spread in percentage points, measured as the difference between the log yield on the ten-year US constant-maturity bond and the log yield on the three-month US Treasury Bill. DEF is the default yield spread in percentage points, measured as the difference between the log yield on Moody's BAA bonds and the log yield on Moody's AAA bonds. VS is the small-stock value-spread, the difference in the log book-to-market ratios of small value and small growth stocks. Bootstrap standard errors (in parentheses) are computed from 2,500 simulated sample realizations.

News cov.	N_{CF}	N_{DR}	N_V	News corr/std	N_{CF}	N_{DR}	N_V
N_{CF}	0.00244	-0.00048	-0.00259	N_{CF}	0.049	-0.108	-0.446
	(0.00094)	(0.00121)	(0.00127)		(0.008)	(0.229)	(0.269)
N_{DR}	-0.00048	0.00802	0.00045	N_{DR}	-0.108	0.090	0.043
	(0.00121)	(0.00285)	(0.00246)		(0.229)	(0.015)	(0.370)
N_V	-0.00259	0.00045	0.01379	N_V	-0.446	0.043	0.117
	(0.00127)	(0.00246)	(0.00482)		(0.269)	(0.370)	(0.028)

Shock correlations	N_{CF}	N_{DR}	N_V	Functions	N_{CF}	N_{DR}	N_V
r_M shock	0.188	-1.640	-0.077	r_M shock	0.958	-0.042	-0.090
	(0.195)	(0.437)	(0.342)		(0.040)	(0.040)	(0.069)
$RVAR$ shock	-0.032	0.773	0.308	$RVAR$ shock	-0.084	-0.084	0.990
	(0.113)	(0.284)	(0.136)		(0.349)	(0.349)	(0.616)
PE shock	0.086	-1.698	-0.146	PE shock	-0.952	-0.952	0.581
	(0.175)	(0.448)	(0.352)		(0.178)	(0.178)	(0.321)
TY shock	0.048	0.281	-0.120	TY shock	0.013	0.013	-0.023
	(0.114)	(0.253)	(0.264)		(0.017)	(0.017)	(0.030)
DEF shock	-0.166	0.530	0.699	DEF shock	-0.069	-0.069	0.392
	(0.143)	(0.376)	(0.248)		(0.044)	(0.044)	(0.080)
VS shock	-0.180	-0.474	0.440	VS shock	-0.187	-0.187	0.304
	(0.125)	(0.269)	(0.264)		(0.132)	(0.132)	(0.232)

Table 5: Cash-flow, Discount-rate, and Variance Betas in the Early Sample

The table shows the estimated cash-flow ($\widehat{\beta}_{CF}$), discount-rate ($\widehat{\beta}_{DR}$), and variance betas ($\widehat{\beta}_{V}$) for the 25 ME- and BE/ME-sorted portfolios (Panel A) and six risk-sorted portfolios (Panel B). "Growth" denotes the lowest BE/ME, "Value" the highest BE/ME, "Small" the lowest ME, and "Large" the highest ME stocks. $\widehat{b}_{\Delta VAR}$ and $\widehat{b}_{r_M}$ are past return-loadings on the weighted sum of changes in the VAR state variables, where the weights are according to λ_V as estimated in Table 4, and on the market-return shock. "Diff." is the difference between the extreme cells. Standard errors [in brackets] are computed from 2,500 simulated sample realizations and are conditional on the estimated news series. Estimates are based on quarterly data for the 1936:3-1963:2 period using weighted least squares. Initial weights of each time-$t+1$ observation are inversely proportional to $EVAR_t$. These weights are then shrunk towards equal weights so that the maximum ratio of actual weights used is bounded by the corresponding historical ratio for the VIX index.

Panel A: 25 ME- and BE/ME-sorted portfolios

$\widehat{\beta}_{CF}$	Growth		2		3		4		Value		Diff	
Small	0.44	[0.13]	0.43	[0.10]	0.41	[0.10]	0.43	[0.10]	0.46	[0.10]	0.02	[0.06]
2	0.32	[0.07]	0.35	[0.09]	0.36	[0.09]	0.40	[0.09]	0.43	[0.10]	0.11	[0.04]
3	0.31	[0.08]	0.30	[0.08]	0.34	[0.09]	0.34	[0.08]	0.47	[0.12]	0.16	[0.05]
4	0.28	[0.07]	0.29	[0.07]	0.33	[0.08]	0.36	[0.08]	0.47	[0.11]	0.19	[0.05]
Large	0.24	[0.07]	0.24	[0.06]	0.28	[0.08]	0.35	[0.09]	0.41	[0.29]	0.17	[0.05]
Diff	-0.20	[0.07]	-0.20	[0.06]	-0.13	[0.05]	-0.08	[0.04]	-0.05	[0.04]		

$\widehat{\beta}_{DR}$	Growth		2		3		4		Value		Diff	
Small	0.89	[0.14]	0.88	[0.15]	0.82	[0.16]	0.81	[0.16]	0.83	[0.15]	-0.06	[0.07]
2	0.69	[0.11]	0.74	[0.13]	0.70	[0.14]	0.73	[0.15]	0.84	[0.12]	0.15	[0.07]
3	0.67	[0.13]	0.63	[0.09]	0.68	[0.11]	0.69	[0.11]	0.80	[0.14]	0.13	[0.08]
4	0.57	[0.07]	0.60	[0.09]	0.63	[0.09]	0.67	[0.12]	0.85	[0.14]	0.28	[0.12]
Large	0.56	[0.08]	0.52	[0.08]	0.52	[0.10]	0.66	[0.13]	0.71	[0.12]	0.15	[0.12]
Diff	-0.33	[0.13]	-0.36	[0.10]	-0.31	[0.15]	-0.15	[0.12]	-0.12	[0.08]		

$\widehat{\beta}_{V}$	Growth		2		3		4		Value		Diff	
Small	-0.72	[0.29]	-0.79	[0.24]	-0.85	[0.26]	-0.82	[0.25]	-0.88	[0.25]	-0.16	[0.13]
2	-0.50	[0.17]	-0.52	[0.22]	-0.60	[0.20]	-0.62	[0.22]	-0.82	[0.25]	-0.32	[0.12]
3	-0.48	[0.20]	-0.38	[0.15]	-0.53	[0.19]	-0.55	[0.19]	-0.85	[0.27]	-0.37	[0.13]
4	-0.21	[0.13]	-0.35	[0.17]	-0.43	[0.18]	-0.58	[0.24]	-0.87	[0.28]	-0.65	[0.19]
Large	-0.22	[0.14]	-0.23	[0.14]	-0.44	[0.21]	-0.67	[0.27]	-0.68	[0.18]	-0.46	[0.16]
Diff	0.50	[0.21]	0.56	[0.14]	0.41	[0.17]	0.16	[0.14]	0.19	[0.12]		

Panel B: 6 risk-sorted portfolios

$\widehat{\beta}_{CF}$	Lo $\widehat{b}_{r_M}$		2		Hi $\widehat{b}_{r_M}$		Diff	
Lo $\widehat{b}_{\Delta VAR}$	0.25	[0.07]	0.36	[0.09]	0.44	[0.11]	0.19	[0.04]
Hi $\widehat{b}_{\Delta VAR}$	0.20	[0.06]	0.26	[0.07]	0.39	[0.10]	0.19	[0.04]
Diff	-0.05	[0.02]	-0.10	[0.03]	-0.05	[0.03]		

$\widehat{\beta}_{DR}$	Lo $\widehat{b}_{r_M}$		2		Hi $\widehat{b}_{r_M}$		Diff	
Lo $\widehat{b}_{\Delta VAR}$	0.56	[0.07]	0.69	[0.10]	0.88	[0.12]	0.32	[0.07]
Hi $\widehat{b}_{\Delta VAR}$	0.42	[0.06]	0.58	[0.08]	0.73	[0.11]	0.31	[0.07]
Diff	-0.14	[0.03]	-0.11	[0.05]	-0.15	[0.05]		

$\widehat{\beta}_{V}$	Lo $\widehat{b}_{r_M}$		2		Hi $\widehat{b}_{r_M}$		Diff	
Lo $\widehat{b}_{\Delta VAR}$	-0.29	[0.14]	-0.51	[0.21]	-0.71	[0.26]	-0.42	[0.13]
Hi $\widehat{b}_{\Delta VAR}$	-0.21	[0.12]	-0.31	[0.16]	-0.52	[0.22]	-0.32	[0.13]
Diff	0.09	[0.07]	0.20	[0.10]	0.19	[0.08]		

Table 6: Cash-flow, Discount-rate, and Variance Betas in the Modern Sample

The table shows the estimated cash-flow ($\widehat{\beta}_{CF}$), discount-rate ($\widehat{\beta}_{DR}$), and variance betas ($\widehat{\beta}_{V}$) for the 25 ME- and BE/ME-sorted portfolios (Panel A), six risk-sorted portfolios (Panel B), and the S&P 100 index straddle portfolio ($STRADDLE$) and the Fama-French factors $RMRF$, SMB, HML, and $DEFRET$ (Panel C). "Growth" denotes the lowest BE/ME, "Value" the highest BE/ME, "Small" the lowest ME, and "Large" the highest ME stocks. $\widehat{b}_{\Delta VAR}$ and $\widehat{b}_{r_M}$ are past return-loadings on the weighted sum of changes in the VAR state variables, where the weights are according to λ_V as estimated in Table 4, and on the market-return shock. "Diff." is the difference between the extreme cells. Standard errors [in brackets] are computed from 2,500 simulated sample realizations and are conditional on the estimated news series. Estimates are based on quarterly data for the 1963:3-2010:4 period in Panels A and B and the 1986:1-2010:4 period in Panel C. Initial weights of each time-$t+1$ observation are inversely proportional to $EVAR_t$. These weights are then shrunk towards equal weights so that the maximum ratio of actual weights used is bounded by the corresponding historical ratio for the VIX index.

Panel A: 25 ME- and BE/ME-sorted portfolios

$\widehat{\beta}_{CF}$	Growth		2		3		4		Value		Diff	
Small	0.22	[0.06]	0.21	[0.05]	0.21	[0.04]	0.20	[0.04]	0.24	[0.05]	0.02	[0.04]
2	0.20	[0.05]	0.19	[0.04]	0.21	[0.04]	0.20	[0.04]	0.23	[0.05]	0.03	[0.04]
3	0.18	[0.05]	0.19	[0.04]	0.18	[0.04]	0.19	[0.04]	0.21	[0.04]	0.03	[0.04]
4	0.17	[0.04]	0.18	[0.04]	0.19	[0.04]	0.18	[0.04]	0.21	[0.05]	0.04	[0.03]
Large	0.11	[0.03]	0.14	[0.03]	0.13	[0.03]	0.14	[0.04]	0.16	[0.04]	0.05	[0.03]
Diff	-0.10	[0.04]	-0.08	[0.03]	-0.08	[0.03]	-0.06	[0.03]	-0.08	[0.03]		

$\widehat{\beta}_{DR}$	Growth		2		3		4		Value		Diff	
Small	1.33	[0.11]	1.09	[0.09]	0.92	[0.08]	0.86	[0.08]	0.89	[0.09]	-0.45	[0.09]
2	1.26	[0.09]	1.00	[0.08]	0.88	[0.07]	0.80	[0.07]	0.82	[0.09]	-0.44	[0.09]
3	1.18	[0.08]	0.92	[0.06]	0.80	[0.07]	0.74	[0.07]	0.76	[0.08]	-0.42	[0.09]
4	1.06	[0.07]	0.89	[0.06]	0.77	[0.06]	0.74	[0.06]	0.79	[0.08]	-0.28	[0.09]
Large	0.89	[0.05]	0.75	[0.05]	0.62	[0.05]	0.61	[0.06]	0.68	[0.06]	-0.21	[0.07]
Diff	-0.45	[0.11]	-0.34	[0.09]	-0.30	[0.07]	-0.24	[0.07]	-0.21	[0.09]		

$\widehat{\beta}_{V}$	Growth		2		3		4		Value		Diff	
Small	0.61	[0.32]	0.37	[0.26]	0.24	[0.24]	0.19	[0.22]	0.02	[0.32]	-0.59	[0.12]
2	0.69	[0.29]	0.42	[0.26]	0.24	[0.23]	0.19	[0.25]	0.08	[0.27]	-0.60	[0.12]
3	0.67	[0.28]	0.36	[0.24]	0.27	[0.22]	0.13	[0.25]	0.15	[0.19]	-0.52	[0.14]
4	0.63	[0.25]	0.36	[0.23]	0.19	[0.26]	0.17	[0.27]	0.11	[0.27]	-0.53	[0.12]
Large	0.47	[0.22]	0.36	[0.17]	0.18	[0.19]	0.12	[0.25]	0.14	[0.21]	-0.33	[0.09]
Diff	-0.14	[0.14]	-0.01	[0.13]	-0.06	[0.09]	-0.07	[0.09]	0.12	[0.14]		

Panel B: 6 risk-sorted portfolios

$\widehat{\beta}_{CF}$	Lo $\widehat{b}_{r_M}$		2		Hi $\widehat{b}_{r_M}$		Diff	
Lo $\widehat{b}_{\Delta VAR}$	0.14	[0.03]	0.17	[0.04]	0.21	[0.06]	0.07	[0.04]
Hi $\widehat{b}_{\Delta VAR}$	0.11	[0.03]	0.15	[0.04]	0.18	[0.05]	0.08	[0.04]
Diff	-0.03	[0.02]	-0.01	[0.02]	-0.02	[0.02]		

$\widehat{\beta}_{DR}$	Lo $\widehat{b}_{r_M}$		2		Hi $\widehat{b}_{r_M}$		Diff	
Lo $\widehat{b}_{\Delta VAR}$	0.67	[0.05]	0.81	[0.06]	1.20	[0.10]	0.54	[0.10]
Hi $\widehat{b}_{\Delta VAR}$	0.64	[0.06]	0.90	[0.05]	1.23	[0.09]	0.59	[0.11]
Diff	-0.02	[0.08]	0.09	[0.07]	0.03	[0.06]		

$\widehat{\beta}_{V}$	Lo $\widehat{b}_{r_M}$		2		Hi $\widehat{b}_{r_M}$		Diff	
Lo $\widehat{b}_{\Delta VAR}$	0.22	[0.21]	0.25	[0.26]	0.47	[0.33]	0.25	[0.16]
Hi $\widehat{b}_{\Delta VAR}$	0.33	[0.17]	0.48	[0.20]	0.60	[0.27]	0.28	[0.14]
Diff	0.10	[0.07]	0.23	[0.09]	0.14	[0.09]		

Panel C: Option, bond, and equity portfolios

	$STRADDLE$		$RMRF$		SMB		HML		$DEFRET$	
$\widehat{\beta}_{CF}$	-0.14	[0.22]	0.17	[0.05]	0.04	[0.02]	0.02	[0.04]	0.06	[0.03]
$\widehat{\beta}_{DR}$	-1.57	[0.47]	0.81	[0.06]	0.19	[0.05]	-0.27	[0.10]	0.20	[0.06]
$\widehat{\beta}_{V}$	1.27	[0.84]	-0.12	[0.34]	-0.02	[0.07]	-0.47	[0.13]	-0.31	[0.26]

Table 7: Asset Pricing Tests for the Early Sample

The table shows the premia estimated from the 1936:3-1963:2 sample for the CAPM, the 2-beta ICAPM, the 3-beta volatility ICAPM, a factor model where only the $\widehat{\beta}_{DR}$ premium is restricted, and an unrestricted factor model. The test assets are the 25 ME- and BE/ME-sorted portfolios. The first column per model constrains the zero-beta rate (R_{zb}) to equal the risk-free rate (R_{rf}) while the second column allows R_{zb} to be a free parameter. Estimates are from a cross-sectional regression of average simple excess test-asset returns (quarterly in fractions) on an intercept and estimated cash-flow ($\widehat{\beta}_{CF}$), discount-rate ($\widehat{\beta}_{DR}$), and variance betas ($\widehat{\beta}_{V}$). Standard errors and critical values [A] are conditional on the estimated news series and (B) incorporate full estimation uncertainty of the news terms. The test rejects if the pricing error is higher than the listed 5 percent critical value.

Parameter	CAPM		2-beta ICAPM		3-beta ICAPM		Constrained		Unrestricted	
$R_{zb} - R_f$ (g_0)	0	-0.006	0	-0.004	0	-0.001	0	0.016	0	0.027
% per annum	0%	-2.41%	0%	-1.49%	0%	-0.23%	0%	6.26%	0%	10.72%
Std. err. A	N/A	[0.017]	N/A	[0.015]	N/A	[0.013]	N/A	[0.017]	N/A	[0.021]
Std. err. B	N/A	(0.017)	N/A	(0.016)	N/A	(0.015)	N/A	(0.017)	N/A	(0.019)
$\widehat{\beta}_{CF}$ prem. (g_1)	0.044	0.049	0.097	0.107	0.080	0.081	0.064	-0.018	0.116	0.053
% per annum	17.50%	19.69%	38.70%	42.67%	32.05%	32.42%	25.59%	-6.99%	46.57%	21.00%
Std. err. A	[0.017]	[0.026]	[0.042]	[0.064]	[0.032]	[0.044]	[0.087]	[0.121]	[0.136]	[0.141]
Std. err. B	(0.016)	(0.026)	(0.127)	(0.106)	(0.072)	(0.083)	(0.103)	(0.119)	(0.143)	(0.146)
$\widehat{\beta}_{DR}$ prem. (g_2)	0.044	0.049	0.017	0.017	0.017	0.017	0.017	0.017	-0.005	-0.037
% per annum	17.50%	19.69%	6.64%	6.64%	6.64%	6.64%	6.64%	6.64%	-2.11%	-14.72%
Std. err. A	[0.017]	[0.026]	[0.006]	[0.006]	[0.006]	[0.006]	[0.006]	[0.006]	[0.068]	[0.082]
Std. err. B	(0.016)	(0.026)	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)	(0.095)	(0.102)
$\widehat{\beta}_{V}$ prem. (g_3)					-0.010	-0.011	-0.020	-0.044	-0.014	-0.046
% per annum					-4.05%	-4.19%	-7.91%	-17.49%	-5.49%	-18.42%
Std. err. A					[0.004]	[0.007]	[0.044]	[0.049]	[0.045]	[0.060]
Std. err. B					(0.042)	(0.038)	(0.158)	(0.150)	(0.218)	(0.239)
$\widehat{R^2}$	51.99%	52.61%	57.02%	57.35%	59.06%	59.07%	59.70%	61.56%	60.34%	64.48%
Pricing error	0.027	0.024	0.022	0.021	0.020	0.019	0.019	0.019	0.018	0.020
5% crit. val. A	[0.074]	[0.034]	[0.066]	[0.040]	[0.070]	[0.044]	[0.044]	[0.034]	[0.042]	[0.038]
5% crit. val. B	(0.076)	(0.034)	(0.088)	(0.045)	(0.113)	(0.051)	(0.047)	(0.036)	(0.043)	(0.042)
Implied γ	N/A	N/A	5.83	6.42	4.82	4.88	N/A	N/A	N/A	N/A
Implied ω	N/A	N/A	N/A	N/A	1.22	1.26	N/A	N/A	N/A	N/A

Table 8: Asset Pricing Tests for the Modern Sample

The table shows the premia estimated from the 1963:3-2010:4 sample for the CAPM, the 2-beta ICAPM, the 3-beta volatility ICAPM, a factor model where only the $\widehat{\beta}_{DR}$ premium is restricted, and an unrestricted factor model. The test assets are the 25 ME- and BE/ME-sorted portfolios. The first column per model constrains the zero-beta rate (R_{zb}) to equal the risk-free rate (R_{rf}) while the second column allows R_{zb} to be a free parameter. Estimates are from a cross-sectional regression of average simple excess test-asset returns (quarterly in fractions) on an intercept and estimated cash-flow ($\widehat{\beta}_{CF}$), discount-rate ($\widehat{\beta}_{DR}$), and variance betas ($\widehat{\beta}_{V}$). Standard errors and critical values [A] are conditional on the estimated news series and (B) incorporate full estimation uncertainty of the news terms. The test rejects if the pricing error is higher than the listed 5 percent critical value.

Parameter	CAPM		2-beta ICAPM		3-beta ICAPM		Constrained		Unrestricted	
$R_{zb} - R_f$ (g_0)	0	0.028	0	-0.013	0	0.011	0	0.001	0	0.003
% per annum	0%	11.06%	0%	-5.05%	0%	4.55%	0%	0.56%	0%	1.02%
Std. err. A	N/A	[0.015]	N/A	[0.013]	N/A	[0.012]	N/A	[0.014]	N/A	[0.013]
Std. err. B	N/A	(0.015)	N/A	(0.018)	N/A	(0.015)	N/A	(0.015)	N/A	(0.016)
$\widehat{\beta}_{CF}$ prem. (g_1)	0.020	-0.005	0.086	0.152	0.056	0.070	0.126	0.119	0.216	0.208
% per annum	8.08%	-1.93%	34.44%	60.76%	22.41%	27.86%	50.30%	47.64%	86.42%	83.22%
Std. err. A	[0.009]	[0.018]	[0.042]	[0.065]	[0.028]	[0.023]	[0.059]	[0.076]	[0.096]	[0.102]
Std. err. B	(0.009)	(0.018)	(0.082)	(0.107)	(0.039)	(0.056)	(0.117)	(0.123)	(0.130)	(0.132)
$\widehat{\beta}_{DR}$ prem. (g_2)	0.020	-0.005	0.008	0.008	0.008	0.008	0.008	0.008	-0.020	-0.021
% per annum	8.08%	-1.93%	3.05%	3.05%	3.05%	3.05%	3.05%	3.05%	-7.89%	-8.39%
Std. err. A	[0.009]	[0.018]	[0.002]	[0.002]	[0.002]	[0.002]	[0.002]	[0.002]	[0.025]	[0.026]
Std. err. B	(0.009)	(0.018)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.055)	(0.058)
$\widehat{\beta}_{V}$ prem. (g_3)					-0.015	-0.029	-0.026	-0.027	-0.002	-0.002
% per annum					-6.14%	-11.75%	-10.57%	-10.75%	-0.71%	-0.60%
Std. err. A					[0.016]	[0.015]	[0.020]	[0.021]	[0.028]	[0.028]
Std. err. B					(0.027)	(0.038)	(0.107)	(0.107)	(0.126)	(0.125)
$\widehat{R}^2$	-38.86%	5.60%	22.83%	30.41%	-118.02%	72.02%	76.00%	76.08%	78.47%	78.72%
Pricing error	0.122	0.116	0.063	0.052	0.242	0.037	0.025	0.025	0.023	0.024
5% crit. val. A	[0.060]	[0.038]	[0.057]	[0.048]	[0.500]	[0.116]	[0.068]	[0.050]	[0.052]	[0.036]
5% crit. val. B	(0.057)	(0.039)	(0.088)	(0.063)	(0.480)	(0.140)	(0.108)	(0.085)	(0.069)	(0.053)
Implied γ	N/A	N/A	11.28	19.90	7.34	9.13	N/A	N/A	N/A	N/A
Implied ω	N/A	N/A	N/A	N/A	4.02	7.69	N/A	N/A	N/A	N/A

Table 9: Asset Pricing Tests for the Early Sample: Inclusion of Risk-sorted Portfolios

The table shows the premia estimated from the 1936:3-1963:2 sample for the CAPM, the 2-beta ICAPM, the 3-beta volatility ICAPM, a factor model where only the $\widehat{\beta}_{DR}$ premium is restricted, and an unrestricted factor model. The test assets are six ME- and BE/ME-sorted portfolios and six risk-sorted portfolios. The first column per model constrains the zero-beta rate (R_{zb}) to equal the risk-free rate (R_{rf}) while the second column allows R_{zb} to be a free parameter. Estimates are from a cross-sectional regression of average simple excess test-asset returns (quarterly in fractions) on an intercept and estimated cash-flow ($\widehat{\beta}_{CF}$), discount-rate ($\widehat{\beta}_{DR}$), and variance betas ($\widehat{\beta}_{V}$). Standard errors and critical values [A] are conditional on the estimated news series and (B) incorporate full estimation uncertainty of the news terms. The test rejects if the pricing error is higher than the listed 5 percent critical value.

Parameter	CAPM		2-beta ICAPM		3-beta ICAPM		Constrained		Unrestricted	
$R_{zb} - R_f$ (g_0)	0	0.000	0	0.000	0	0.000	0	0.029	0	0.035
% per annum	0%	0.13%	0%	0.15%	0%	-0.05%	0%	11.55%	0%	13.98%
Std. err. A	N/A	[0.016]	0	[0.015]	N/A	[0.013]	0	[0.017]	N/A	[0.018]
Std. err. B	N/A	(0.016)	0	(0.015)	N/A	(0.013)	0	(0.016)	N/A	(0.018)
$\widehat{\beta}_{CF}$ prem. (g_1)	0.040	0.040	0.087	0.086	0.074	0.075	0.027	-0.129	0.044	-0.034
% per annum	16.10%	15.98%	34.71%	34.29%	29.76%	29.85%	10.67%	-51.76%	17.64%	-13.40%
Std. err. A	[0.016]	[0.025]	[0.041]	[0.062]	[0.031]	[0.042]	[0.112]	[0.143]	[0.168]	[0.182]
Std. err. B	(0.016)	(0.025)	(0.113)	(0.100)	(0.071)	(0.074)	(0.127)	(0.143)	(0.182)	(0.200)
$\widehat{\beta}_{DR}$ prem. (g_2)	0.040	0.040	0.017	0.017	0.017	0.017	0.017	0.017	0.009	-0.038
% per annum	16.10%	15.98%	6.64%	6.64%	6.64%	6.64%	6.64%	6.64%	3.68%	-15.22%
Std. err. A	[0.016]	[0.025]	[0.006]	[0.006]	[0.006]	[0.006]	[0.006]	[0.006]	[0.094]	[0.106]
Std. err. B	(0.016)	(0.025)	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)	(0.125)	(0.144)
$\widehat{\beta}_{VAR}$ prem. (g_3)					-0.008	-0.008	-0.038	-0.088	-0.036	-0.084
% per annum					-3.28%	-3.31%	-15.22%	-35.00%	-14.45%	-33.47%
Std. err. A					[0.004]	[0.006]	[0.060]	[0.070]	[0.069]	[0.081]
Std. err. B					(0.041)	(0.031)	(0.197)	(0.203)	(0.317)	(0.356)
$\widehat{R^2}$	53.33%	53.33%	57.44%	57.45%	60.78%	60.78%	66.18%	77.87%	66.24%	80.79%
characteristics	47.02%	46.96%	53.21%	53.09%	57.62%	57.68%	66.84%	74.85%	67.35%	78.96%
risk-sorted	64.79%	65.21%	58.28%	59.08%	56.87%	56.50%	41.29%	83.50%	38.58%	80.73%
Pricing error	0.012	0.012	0.011	0.011	0.010	0.010	0.014	0.007	0.014	0.009
5% crit. val. A	[0.044]	[0.017]	[0.038]	[0.019]	[0.046]	[0.024]	[0.031]	[0.024]	[0.022]	[0.021]
5% crit. val. B	(0.044)	(0.018)	(0.048)	(0.024)	(0.064)	(0.030)	(0.032)	(0.024)	(0.025)	(0.026)
Implied γ	N/A	N/A	5.22	5.16	4.48	4.49	N/A	N/A	N/A	N/A
Implied ω	N/A	N/A	N/A	N/A	0.99	1.00	N/A	N/A	N/A	N/A

Table 10: Asset Pricing Tests for the Modern Sample: Inclusion of Risk-sorted Portfolios

The table shows the premia estimated from the 1963:3-2010:4 sample for the CAPM, the 2-beta ICAPM, the 3-beta volatility ICAPM, a factor model where only the $\widehat{\beta}_{DR}$ premium is restricted, and an unrestricted factor model. The test assets are six ME- and BE/ME-sorted portfolios and six risk-sorted portfolios. The first column per model constrains the zero-beta rate (R_{zb}) to equal the risk-free rate (R_{rf}) while the second column allows R_{zb} to be a free parameter. Estimates are from a cross-sectional regression of average simple excess test-asset returns (quarterly in fractions) on an intercept and estimated cash-flow ($\widehat{\beta}_{CF}$), discount-rate ($\widehat{\beta}_{DR}$), and variance betas ($\widehat{\beta}_{V}$). Standard errors and critical values [A] are conditional on the estimated news series and (B) incorporate full estimation uncertainty of the news terms. The test rejects if the pricing error is higher than the listed 5 percent critical value.

Parameter	CAPM		2-beta ICAPM		3-beta ICAPM		Constrained		Unrestricted	
$R_{zb} - R_f$ (g_0)	0	0.015	0	-0.004	0	0.011	0	0.007	0	0.008
% per annum	0%	5.87%	0%	-1.61%	0%	4.25%	0%	2.64%	0%	3.06%
Std. err. A	N/A	[0.010]	0	[0.010]	N/A	[0.010]	0	[0.012]	N/A	[0.011]
Std. err. B	N/A	(0.010)	0	(0.012)	N/A	(0.011)	0	(0.013)	N/A	(0.013)
$\widehat{\beta}_{CF}$ prem. (g_1)	0.017	0.004	0.073	0.095	0.053	0.072	0.125	0.094	0.066	-0.003
% per annum	6.75%	1.56%	29.03%	38.07%	21.14%	28.63%	49.94%	37.50%	26.23%	-1.37%
Std. err. A	[0.009]	[0.014]	[0.044]	[0.061]	[0.030]	[0.032]	[0.069]	[0.083]	[0.140]	[0.153]
Std. err. B	(0.009)	(0.014)	(0.079)	(0.095)	(0.039)	(0.052)	(0.125)	(0.139)	(0.148)	(0.161)
$\widehat{\beta}_{DR}$ prem. (g_2)	0.017	0.004	0.008	0.008	0.008	0.008	0.008	0.008	0.025	0.034
% per annum	6.75%	1.56%	3.05%	3.05%	3.05%	3.05%	3.05%	3.05%	9.80%	13.57%
Std. err. A	[0.009]	[0.014]	[0.002]	[0.002]	[0.002]	[0.002]	[0.002]	[0.002]	[0.035]	[0.036]
Std. err. B	(0.009)	(0.014)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.055)	(0.061)
$\widehat{\beta}_{VAR}$ prem. (g_3)					-0.013	-0.032	-0.028	-0.031	-0.044	-0.056
% per annum					-5.19%	-12.79%	-11.16%	-12.44%	-17.44%	-22.43%
Std. err. A					[0.019]	[0.019]	[0.025]	[0.026]	[0.046]	[0.047]
Std. err. B					(0.026)	(0.033)	(0.111)	(0.120)	(0.140)	(0.157)
$\widehat{R}^2$	-8.47%	9.89%	26.69%	28.08%	-27.31%	75.42%	73.59%	76.71%	74.24%	78.21%
characteristics	-4.07%	15.03%	31.03%	33.39%	12.34%	84.02%	76.65%	80.05%	78.15%	81.52%
risk-sorted	13.52%	23.89%	37.65%	35.78%	-101.02%	68.83%	73.21%	74.52%	71.21%	75.63%
Pricing error	0.050	0.048	0.031	0.031	0.080	0.016	0.016	0.014	0.016	0.014
5% crit. val. A	[0.037]	[0.023]	[0.038]	[0.025]	[0.265]	[0.083]	[0.042]	[0.029]	[0.028]	[0.018]
5% crit. val. B	(0.037)	(0.023)	(0.047)	(0.027)	(0.257)	(0.086)	(0.057)	(0.044)	(0.034)	(0.023)
Implied γ	N/A	N/A	9.51	12.47	6.92	9.38	N/A	N/A	N/A	N/A
Implied ω	N/A	N/A	N/A	N/A	3.40	8.38	N/A	N/A	N/A	N/A

Table 11: Asset Pricing Tests for the Equity and Option Sample

The table shows the premia estimated from the 1986:1-2010:4 sample for the CAPM, the 2-beta ICAPM, the 3-beta volatility ICAPM, a factor model where only the $\widehat{\beta}_{DR}$ premium is restricted, and an unrestricted factor model. The test assets are the three equity factors and the risky bond factor of Fama and French (1993) and the S&P 100 index straddle return from Coval and Shumway (2001). The first column per model constrains the zero-beta rate (R_{zb}) to equal the risk-free rate (R_{rf}) while the second column allows R_{zb} to be a free parameter. Estimates are from a cross-sectional regression of average simple excess test-asset returns (quarterly in fractions) on an intercept and estimated cash-flow ($\widehat{\beta}_{CF}$), discount-rate ($\widehat{\beta}_{DR}$), and variance betas ($\widehat{\beta}_{V}$). Standard errors and critical values [A] are conditional on the estimated news series and (B) incorporate full estimation uncertainty of the news terms. The test rejects if the pricing error is higher than the listed 5 percent critical value.

Parameter	CAPM		2-beta ICAPM		3-beta ICAPM		Constrained		Unrestricted	
R_{zb} less R_f (g_0)	0	-0.031	0	-0.064	0	-0.034	0	-0.040	0	-0.013
% per annum	0%	-12.47%	0%	-25.39%	0%	-13.58%	0%	-15.88%	0%	-5.24%
Std. err. A	N/A	[0.011]	0	[0.037]	N/A	[0.017]	0	[0.039]	N/A	[0.044]
Std. err. B	N/A	(0.011)	0	(0.027)	N/A	(0.014)	0	(0.023)	N/A	(0.032)
$\widehat{\beta}_{CF}$ premium (g_1)	0.103	0.099	0.626	0.802	0.110	0.109	-0.010	0.261	-0.574	-0.324
% per annum	41.27%	39.76%	250.29%	320.76%	44.05%	43.65%	-4.15%	104.20%	-229.60%	-129.60%
Std. err. A	[0.026]	[0.026]	[0.428]	[0.593]	[0.031]	[0.029]	[0.411]	[0.670]	[0.543]	[1.188]
Std. err. B	(0.028)	(0.029)	(0.450)	(0.546)	(0.089)	(0.091)	(0.363)	(0.499)	(0.615)	(0.882)
$\widehat{\beta}_{DR}$ premium (g_2)	0.103	0.004	0.008	0.008	0.008	0.008	0.008	0.008	0.025	0.034
% per annum	41.27%	39.76%	3.17%	3.17%	3.17%	3.17%	3.17%	3.17%	48.42%	35.53%
Std. err. A	[0.026]	[0.014]	[0.001]	[0.001]	[0.002]	[0.002]	[0.001]	[0.001]	[0.035]	[0.036]
Std. err. B	(0.028)	(0.014)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.055)	(0.061)
$\widehat{\beta}_{VAR}$ premium (g_3)					-0.137	-0.132	-0.151	-0.113	-0.098	-0.101
% per annum					-54.82%	-52.66%	-60.46%	-45.19%	-39.22%	-40.23%
Std. err. A					[0.101]	[0.104]	[0.140]	[0.206]	[0.121]	[0.160]
Std. err. B					(0.133)	(0.138)	(0.253)	(0.280)	(0.377)	(0.445)
$\widehat{R^2}$	78.25%	86.29%	50.39%	81.36%	88.15%	97.70%	88.81%	98.48%	99.30%	99.50%
Pricing error	1.742	1.500	2.105	2.010	1.980	0.483	1.902	0.373	0.184	0.134
5% critic. val. A	[0.734]	[0.702]	[2.488]	[3.277]	[3.280]	[3.900]	[3.911]	[2.592]	[1.085]	[0.307]
5% critic. val. B	(0.753)	(0.776)	(3.099)	(3.527)	(3.455)	(3.299)	(2.995)	(1.249)	(1.256)	(0.395)
Implied γ	N/A	N/A	78.92	101.14	13.89	13.76	N/A	N/A	N/A	N/A
Implied ω	N/A	N/A	N/A	N/A	34.57	33.21	N/A	N/A	N/A	N/A

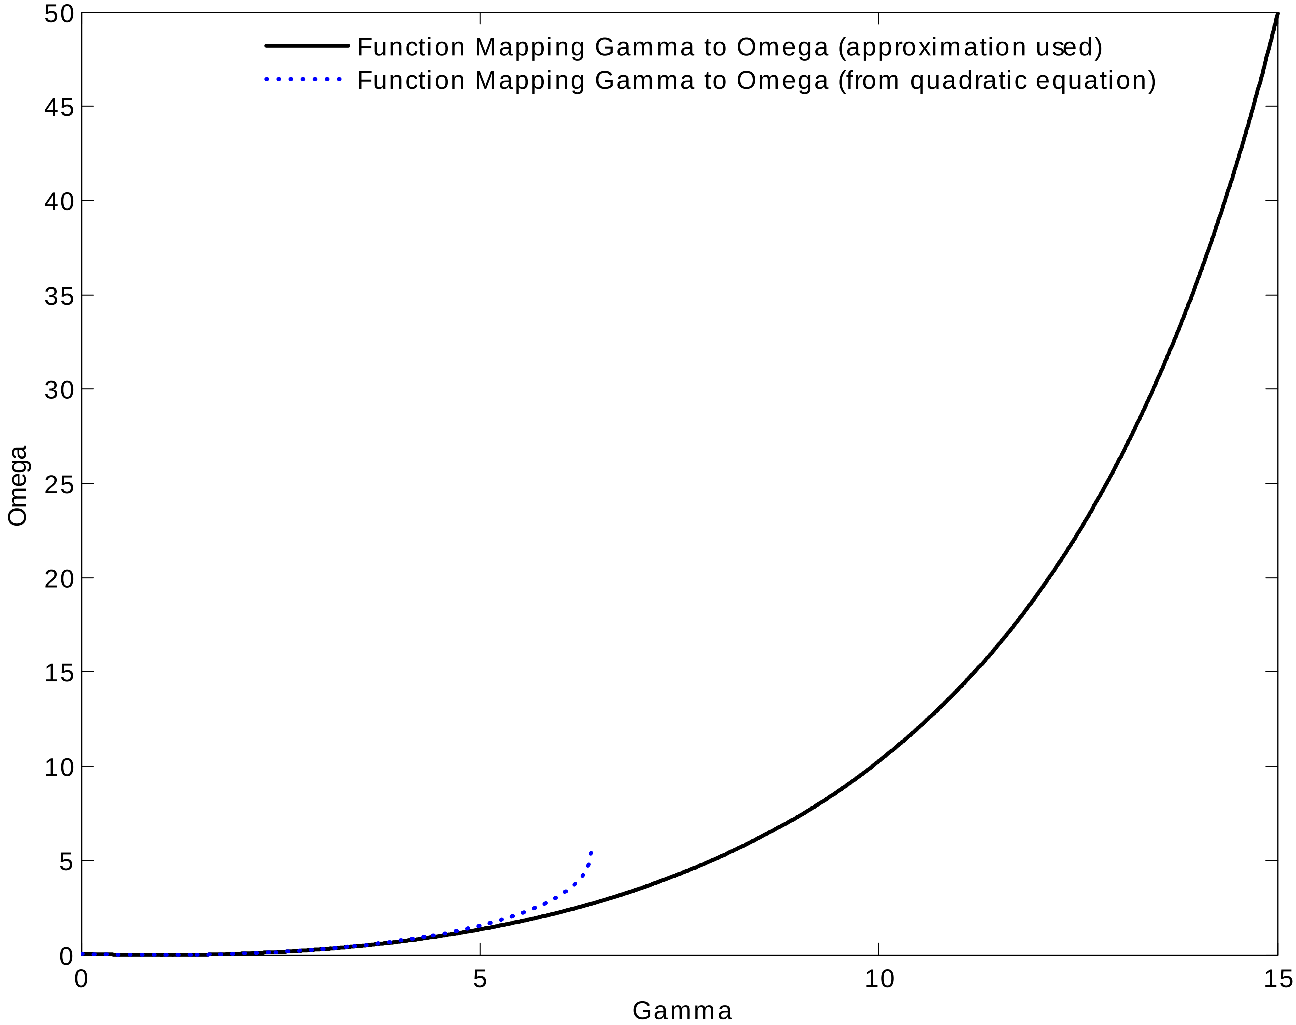

Figure 1: This figure graphs the approximate relation between the parameter γ and the parameter ω described by equation (24) as well as the quadratic solution for ω described in equation (23). These functions depend on the loglinearization parameter ρ, set to 0.95 per year and the empirically estimated VAR parameters of Table 2. γ is the investor's risk aversion while ω is the sensitivity of news about risk, N_{RISK}, to news about market variance, N_V.

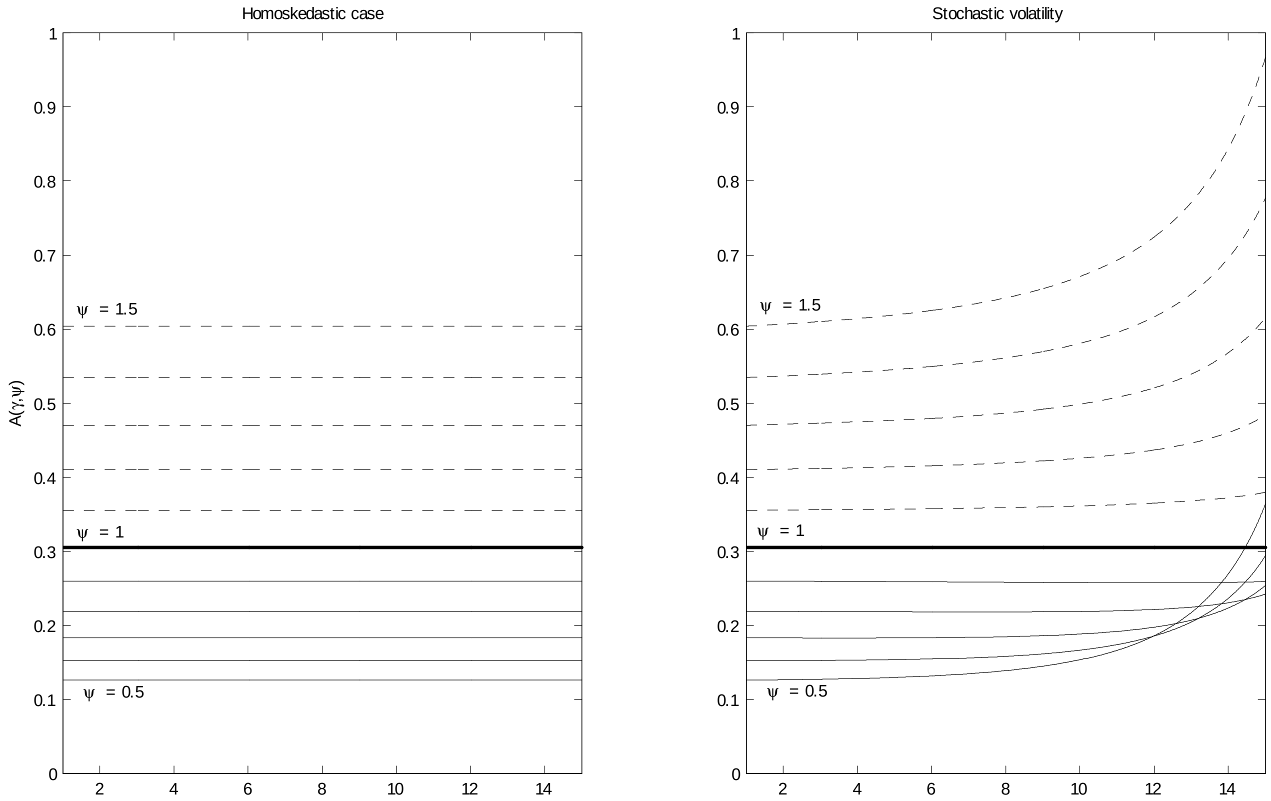

Figure 2: This figure plots plots the coefficient $A(\gamma, \psi)$ relating the conditional variance of consumption growth to σ_t^2 for different values of γ and ψ for the homoskedastic case (left panel) and for the heteroskedastic case (right panel), where $A(\gamma, \psi)$ is a function of the variances and covariances of the *scaled* residuals u_{t+1}. In each panel, we plot $A(\gamma, \psi)$ as γ varies between 0 and 15, for different values of ψ. Each line corresponds to a different ψ between 0.5 and 1.5.

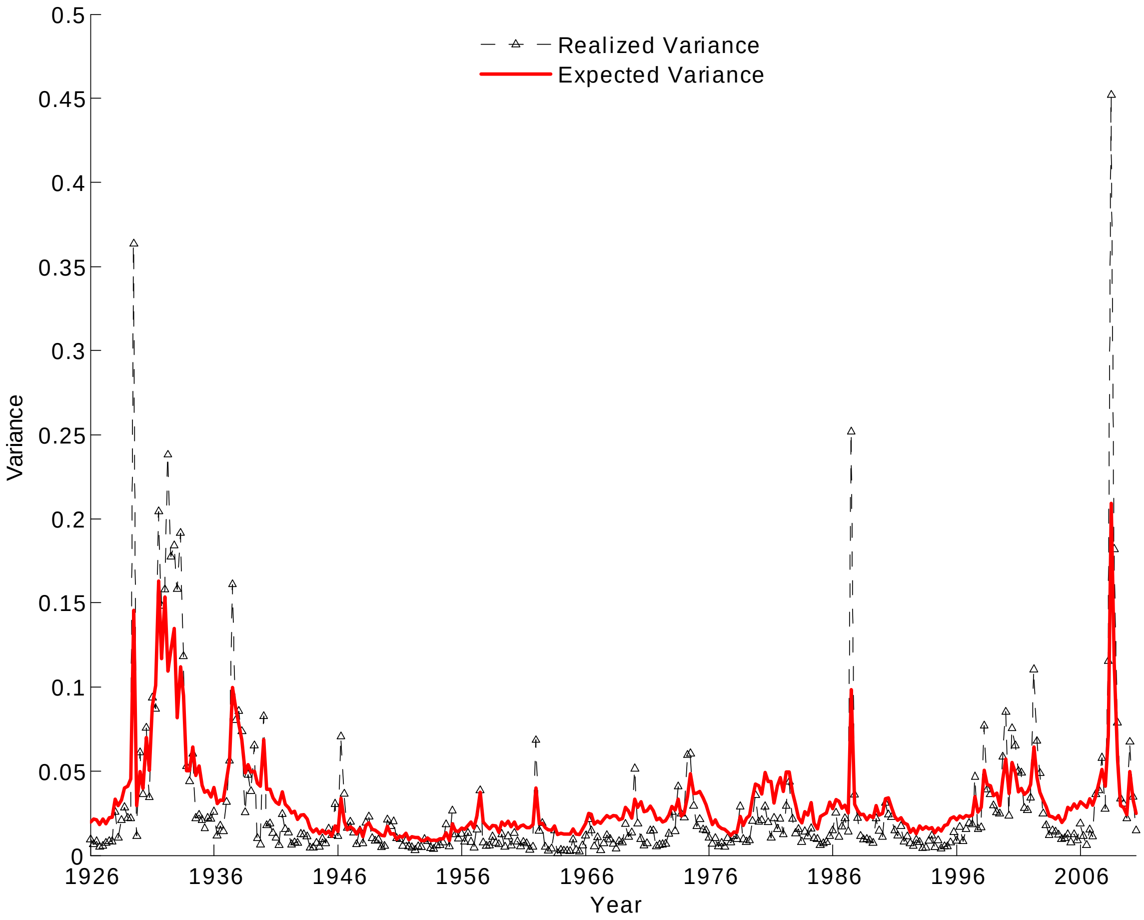

Figure 3: This figure plots quarterly observations of realized within-quarter daily return variance over the sample period 1926:2-2010:4 and the expected variance implied by the model estimated in Table 1 Panel A.

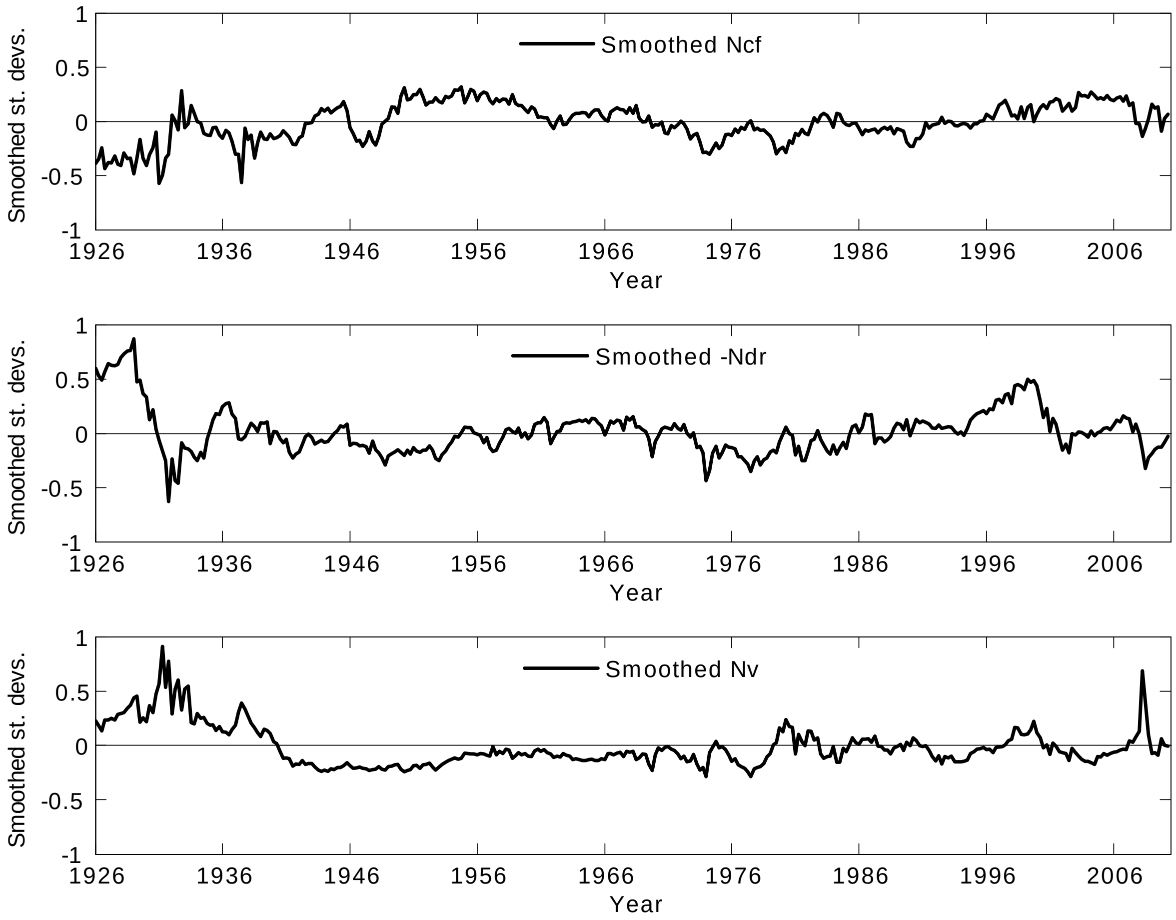

Figure 4: This figure plots normalized cash-flow news, the negative of normalized discount-rate news, and normalized variance news. The series are smoothed with a trailing exponentially-weighted moving average where the decay parameter is set to 0.08 per quarter, and the smoothed news series is generated as $MA_t(N) = 0.08N_t + (1 - 0.08)MA_{t-1}(N)$. This decay parameter implies a half-life of six years. The sample period is 1926:2-2010:4.

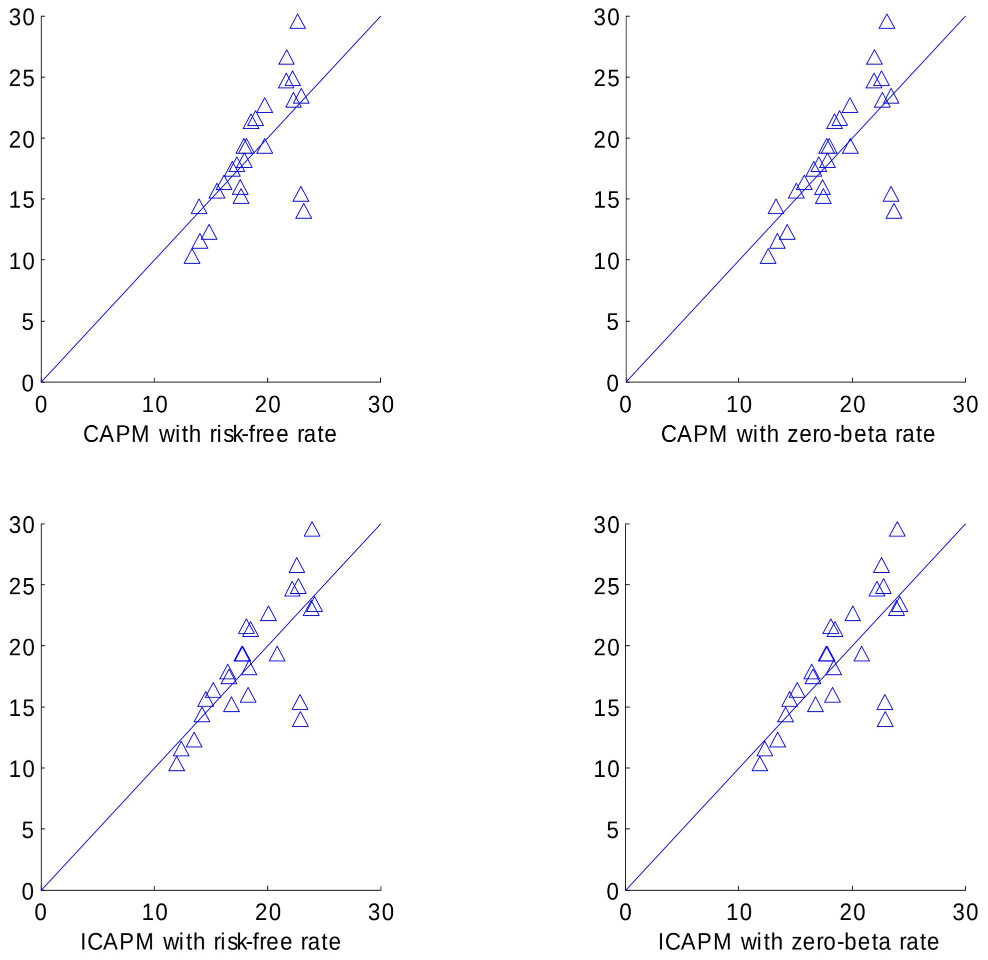

Figure 5: The four diagrams correspond to (clockwise from the top left) the CAPM with a constrained zero-beta rate, the CAPM with an unconstrained zero-beta rate, the three-factor ICAPM with a free zero-beta rate, and the three-factor ICAPM with the zero-beta rate constrained to the risk-freee rate. The horizontal axes correspond to the predicted average excess returns and the vertical axes to the sample average realized excess returns for the 25 ME- and BE/ME-sorted portfolios. The predicted values are from regressions presented in Table 7 for the sample period 1936:3-1963:2.

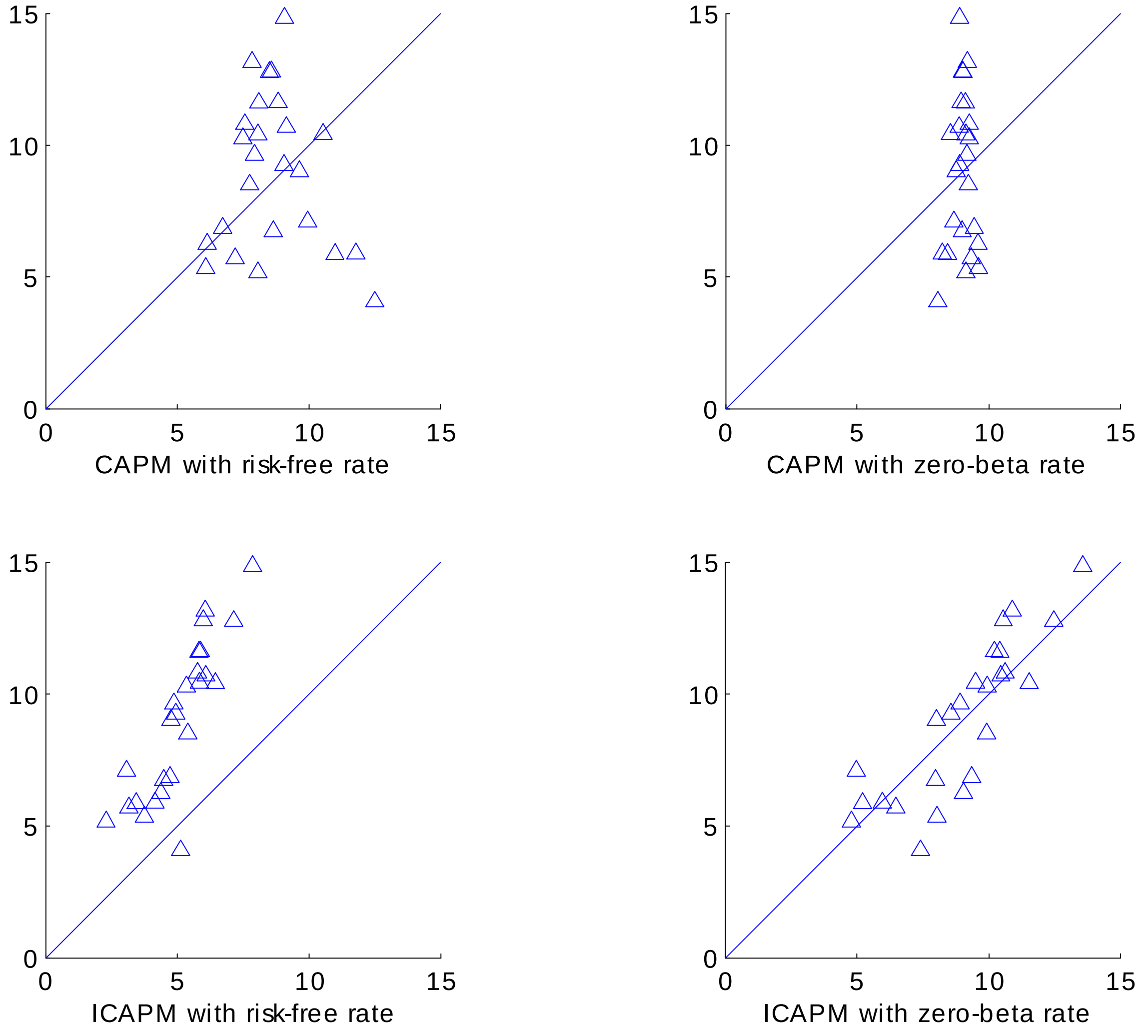

Figure 6: The four diagrams correspond to (clockwise from the top left) the CAPM with a constrained zero-beta rate, the CAPM with an unconstrained zero-beta rate, the three-factor ICAPM with a free zero-beta rate, and the three-factor ICAPM with the zero-beta rate constrained to the risk-freee rate. The horizontal axes correspond to the predicted average excess returns and the vertical axes to the sample average realized excess returns for the 25 ME- and BE/ME-sorted portfolios. The predicted values are from regressions presented in Table 8 for the sample period 1963:3-2010:4.

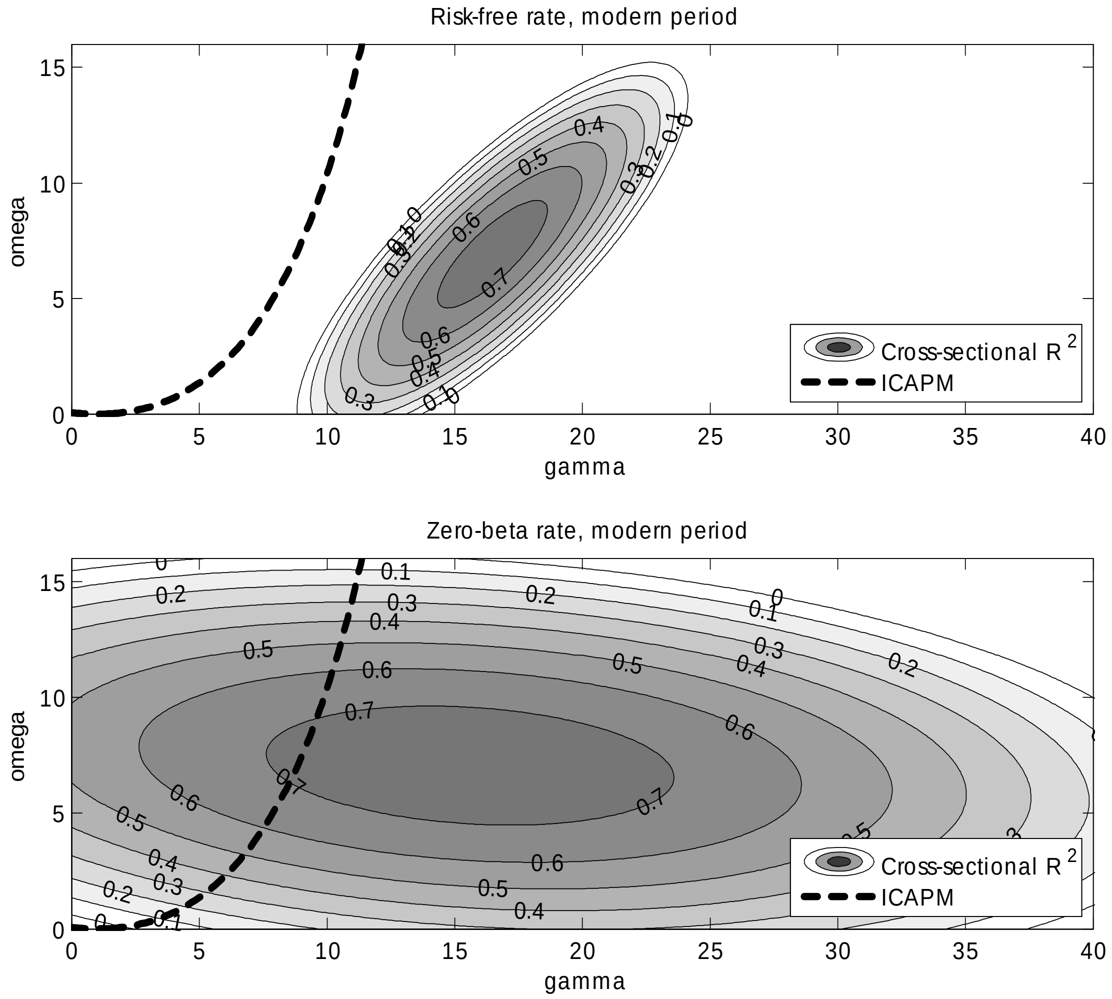

Figure 7: The two contour plots show how the R^2 of the cross-sectional regression explaining the average returns on the 25 size- and book-to-market portfolios varies for different values of γ and ω for the risk-free rate (top panel) and zero-beta rate (bottom panel) three-beta ICAPM model estimated in Table 8 for the sample period 1963:3-2010:4. The two plots also indicate the approximate ICAPM relation between γ and ω described in the equation (24).

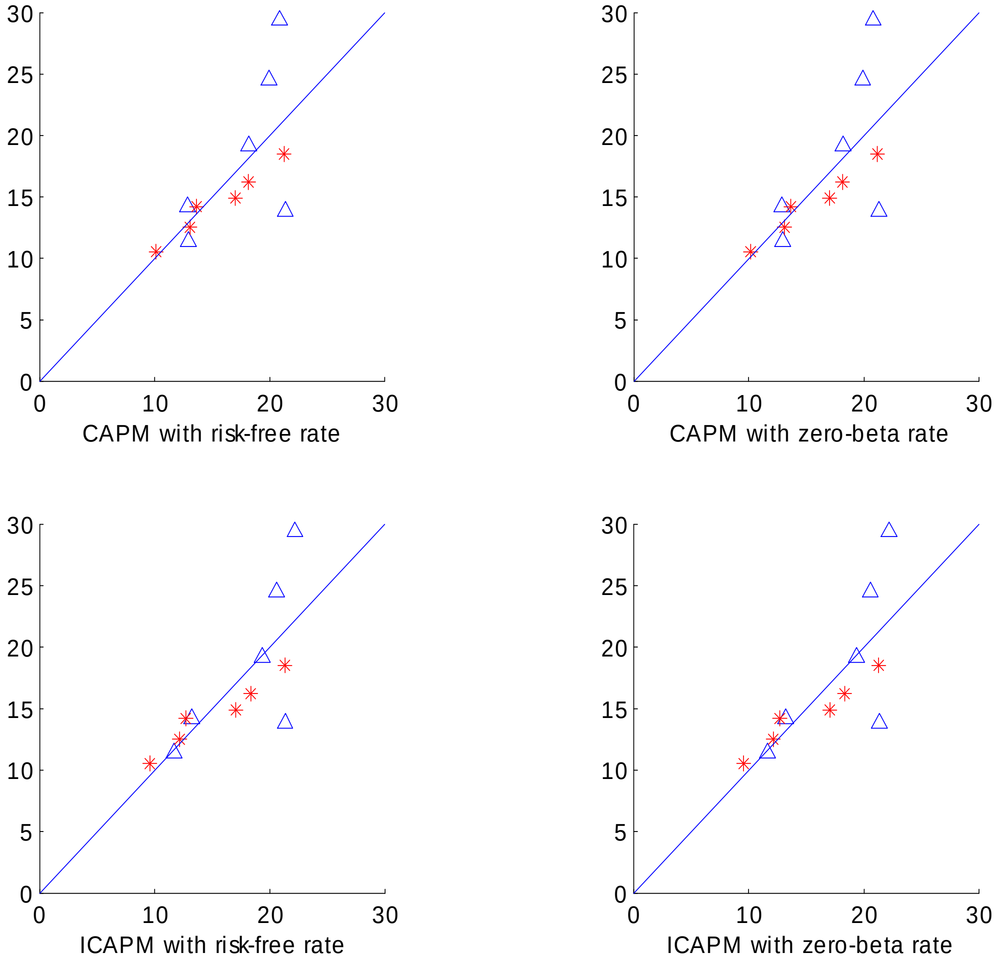

Figure 8: The four diagrams correspond to (clockwise from the top left) the CAPM with a constrained zero-beta rate, the CAPM with an unconstrained zero-beta rate, the three-factor ICAPM with a free zero-beta rate, and the three-factor ICAPM with the zero-beta rate constrained to the risk-freee rate. The horizontal axes correspond to the predicted average excess returns and the vertical axes to the sample average realized excess returns for six ME- and BE/ME-sorted portfolios (denoted by triangles) and six risk-sorted portfolios (denoted by asterisks). The predicted values are from regressions presented in Table 9 for the sample period 1936:3-1963:2.

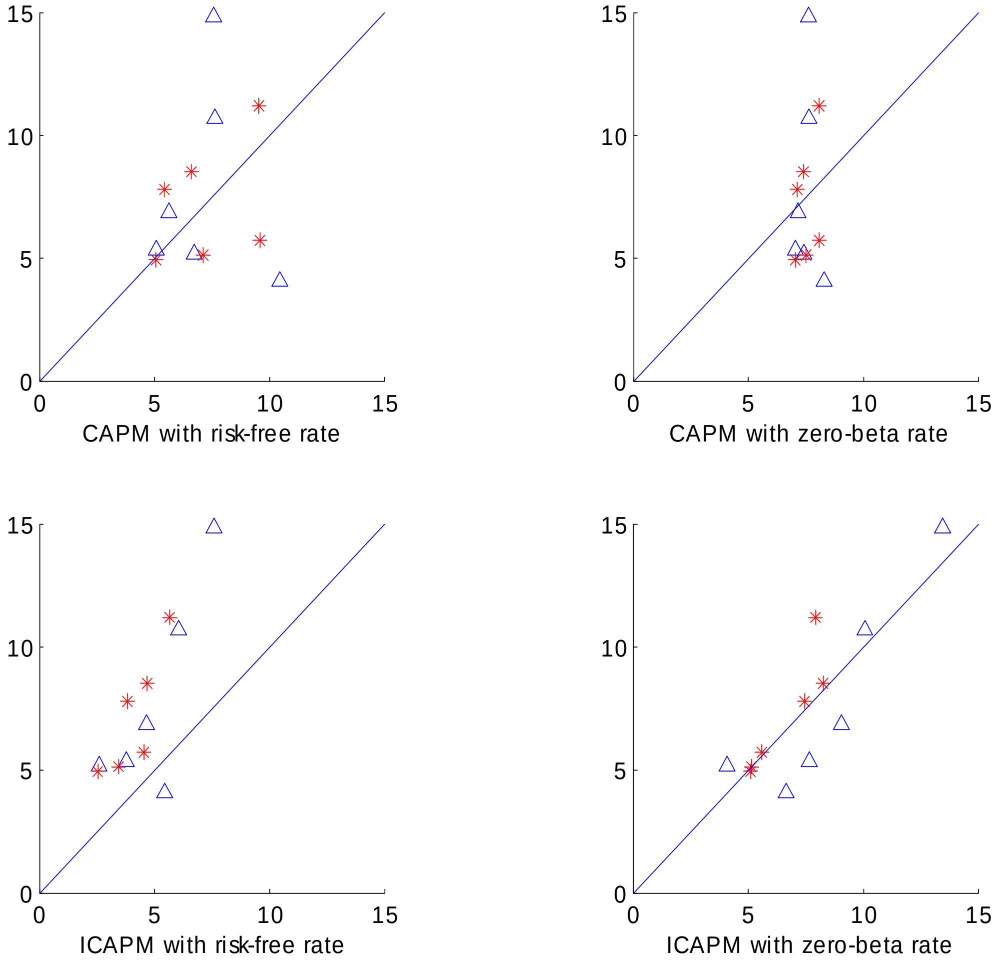

Figure 9: The four diagrams correspond to (clockwise from the top left) the CAPM with a constrained zero-beta rate, the CAPM with an unconstrained zero-beta rate, the three-factor ICAPM with a free zero-beta rate, and the three-factor ICAPM with the zero-beta rate constrained to the risk-freee rate. The horizontal axes correspond to the predicted average excess returns and the vertical axes to the sample average realized excess returns for six ME- and BE/ME-sorted portfolios (denoted by triangles) and six risk-sorted portfolios (denoted by asterisks). The predicted values are from regressions presented in Table 10 for the sample period 1963:3-2010:4.

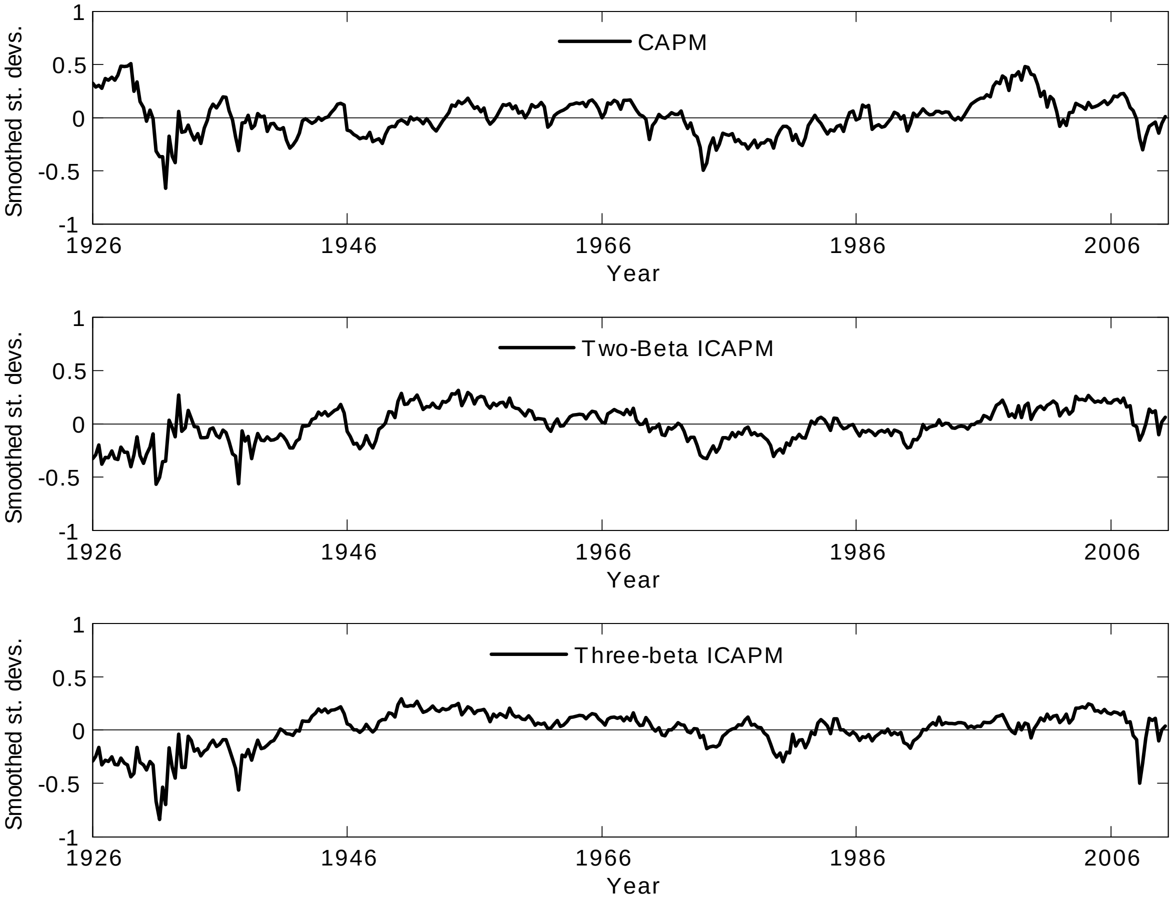

Figure 10: This figure plots the time-series of the smoothed combined shock for the CAPM $(N_{CF}-N_{DR})$, the two-beta ICAPM $(\gamma N_{CF}-N_{DR})$, and the three-beta ICAPM that includes stochastic volatility $(\gamma N_{CF}-N_{DR}-\frac{1}{2}\omega N_V)$ for the unconstrained zero-beta rate specifications estimated in Table 8 for the modern subperiod. The shock is smoothed with a trailing exponentially-weighted moving average. The decay parameter is set to 0.08 per quarter, and the smoothed news series is generated as $MA_t(SDF) = 0.08SDF_t + (1-0.08)MA_{t-1}(N)$. This decay parameter implies a half-life of six years. The sample period is 1926:2-2010:4.