Time-varying predictability of consumption growth, macro-uncertainty, and risk premiums *

Pedro Barroso † Martijn Boons‡ Paul Karehnke §

Abstract

We show that the relation between state variables, such as the t-bill rate and term spread, and consumption growth varies significantly over time. Consistent with an Intertemporal CAPM, we find that state variable risk premiums (in the cross section of individual stocks) vary over time accordingly: Risk premiums increase by 5% (annualized) when a state variable predicts consumption growth strongly relative to its own history. This effect is magnified by time-variation in the variance of the state variables, which we argue to be associated to general macroeconomic uncertainty. Our conditional evidence contributes to recent literature that focuses on the unconditional pricing of state variable risk and finds mixed results.

JEL Classification Codes: G12

Keywords: Conditional asset pricing models, state variable risk premiums, Intertemporal CAPM, macroeconomic uncertainty, time-varying consumption predictability.

*We thank seminar participants at ESSEC, Nova SBE in Lisbon, UNSW in Sydney, and the European Economics Association in Lisbon for helpful comments. This work was funded by National Funds through FCT (Fundao para a Ciencia e Tecnologia, project: UID/ECO/00124/2013) and POR Lisboa (project: LISBOA-01-0145-FEDER-007722).

†The University of New South Wales, E-mail: p.barroso@unsw.edu.au.

‡Nova School of Business and Economics and NETSPAR, E-mail: martijn.boons@novasbe.pt.

§The University of New South Wales, E-mail: p.karehnke@unsw.edu.au.

State variables, such as the t-bill rate and term spread, contain important information about the future performance of assets in the consumption-investment opportunity set. This informational content can change over time, however. A state variable whose shocks contain on average good news about the economy can for prolonged periods convey no news at all or even bad news. In this paper, we analyze the conditional asset pricing implications of this low frequency variation in state variable risk in the cross section of individual stocks. We find that risk premiums for exposure to state variables that predict consumption growth are time-varying in a manner consistent with the time-variation in this predictive relation.[1] That is, when a state variable predicts consumption growth unusually strongly at a certain point in time, we find that its risk premium is also larger than usual. Moreover, we find that the risk premium for a state variable is increasing in the conditional variance of the state variable and we argue that this variance is associated to general macroeconomic uncertainty. In all, we conclude that state variable risk premiums are conditionally consistent with an Intertemporal CAPM (ICAPM).

Our paper follows the advice in Campbell and Cochrane (2000), Cochrane (2005, Ch. 8) and Nagel and Singleton (2011), who argue that conditional tests have more statistical power to distinguish between competing asset pricing models. A recent literature analyzes whether unconditional risk premiums for exposure to state variables line up with macro-finance theory and, in particular, the ICAPM of Merton (1973). This analysis responds to Lewellen and Nagel (2006) and Lewellen et al. (2010), who argue that asset pricing tests should impose theoretical restrictions on risk prices, and Fama (1991), who argues that the state variables considered in the asset pricing test should be motivated by economic hedging arguments. However, as noted in Barbalau, Robotti, and Shanken (2015), the evidence so far is mixed and inconclusive as to whether unconditional risk premiums in the stock market for exposure to state variables are consistent with the ICAPM.

[1]Consumption growth is our main proxy for macroeconomic growth, but we show that our conclusions hold equally for industrial production and GDP growth.

We present four results that hold equally for four empirical multi-factor models that have shown strong explanatory power for the cross section and time series.[2] First, we show that the coefficients in a predictive regression of consumption growth on lagged state variables are time-varying. In a simple two-stage test, controlling for these time-varying coefficients substantially improves consumption growth predictability relative to an unconditional model that includes the same state variables.

Second, we show that these time-varying coefficients predict the risk premiums for exposure to the state variables with a positive sign in the time series. Following recent literature (see, e.g., Gagliardini et al. (2012), Novy-Marx (2015), and Ang et al. (2016)), we estimate these risk premiums in the cross section of individual stocks. We find that state variable risk premiums are larger by about 5% annualized (or, 0.35 in Sharpe ratio) whenever a state variable predicts consumption growth relatively strongly. This effect is statistically significant and economically large, implying that risk premiums change over time about as much as the market risk premium.

Third, we show that risk premiums are increasing (in absolute magnitude) in the conditional variance of the state variables. When the conditional variance of a state variable is high, the effect of a time-varying relation between the state variable and future consumption growth is about twice as strong: we estimate a risk premium of about 5% (-5%) for a state variable that predicts consumption growth positively (negatively) and when variance is high versus about 1% (-1%) when variance is low.

Fourth, we argue that a large share of the variation in the conditional variance of state variables is common and related to general macroeconomic uncertainty, which we measure as the variance of consumption growth conditional on the state variables. Our estimates

[2] Following Boons (2016), our first model includes the dividend yield (DY), the default spread (DS), and the term spread (TS). Our second model substitutes the three-month t-bill rate (RF) for TS, as in, e.g., Petkova (2006). Our third model includes TS, the price-to-earnings ratio (PE), and the value spread (VS) of Campbell and Vuolteenaho (2004). Our fourth model has the Cochrane and Piazzesi (2005) bond market factor (CP) and a term structure level factor (LVL).

are robust in two important dimensions: (i) using alternative proxies for macroeconomic growth and uncertainty based on industrial production and GDP, and (ii) using portfolio-level estimates of state variable risk premiums. We verify in a simulation study that our results are most likely generated in a conditional ICAPM world.

To motivate our evidence theoretically, we propose a conditional sign restriction on the state variable risk premiums that follows from a stochastic discount factor (SDF) that prices systematic economic news and, thus, state variables containing this news. The unconditional implication from this SDF is identical to the ICAPM in that the time series and cross section should be consistent: If a state variable predicts the macroeconomy with a positive sign in the time series, it should capture a positive risk premium in the cross section. In this paper, we allow the SDF-coefficients to be time-varying and proportional to the time t conditional relation between the state variables and the macroeconomy. We test the conditional implication that the time-varying relation between the state variable and the macroeconomy should predict the state variable risk premium with a positive sign in the time series.

To see why this time-variation is important empirically, consider Panel B of Figure 1. Here, we present the conditional relation between consumption growth and the one quarter lagged three-month t-bill rate (denoted RF), estimated with a ten-year backward-looking rolling window regression. We see that the relation between the two series has changed considerably over time.[3] The predictive coefficient is large and negative in the 1970s and 1980s, suggesting that a shock to RF is bad news. In contrast, the coefficient is close to zero over the last decade, suggesting that shocks to RF are not bad news anymore.[4] In line with standard ICAPM intuition, we argue that this variation should be mirrored in the risk

[3]It is perhaps unsurprising that the estimated coefficient from a ten-year rolling window regression is persistent and varies over time, like any near unit root process. However, this variation is not spurious and contributes to the predictability of consumption growth out-of-sample.

[4]Panel D of Figure 1 shows consistent evidence for the term structure level factor. Panel A of Figure 1 shows the opposite pattern for the term spread. Consistent with our results, Harvey (1988) observes that predictive coefficients for consumption growth vary across subsamples, and in an opposite manner for level and slope variables.

premium investors require for exposure to shocks to RF. Stocks co-moving with RF should contain a negative risk premium in the 1970s and 1980s, but not anymore.

As one can think of RF as a proxy for short-term expected inflation, our findings for RF are consistent with the recent reversal in the relation between inflation and the real economy (documented in Bekaert and Wang (2010), Campbell et al. (2014), David and Veronesi (2014), Kang and Pflueger (2015) and Boons et al. (2016)). Similar to us, these authors show that such time-varying inflation risk has an important impact on prices in stock as well as corporate and government bond markets. In particular, Boons et al. (2016) derive a long-run risk model that incorporates this time-variation. More generally, our work contributes to literature that analyzes the conditional asset pricing implications of macroeconomic risks (see the seminal work of Campbell and Cochrane (1999) on habits and of Bansal and Yaron (2004) on long-run risk). We show that time-variation in the relation between a large set of state variables and the macroeconomy has an important impact on equity risk premiums. Natural candidates for factors driving this variation are monetary policy, the business cycle, and the prevalence of supply and demand shocks.

Also consistent with the ICAPM, we argue that time-variation in the conditional variance of a state variable magnifies the time-variation in the state variable risk premium. This contributes to literature on the time-variation in the aggregate equity risk premium that is either due to variance (i.e., the risk-return trade-off studied in French et al. (1987), Duffee (2005) and Ghysels et al. (2005), among many others) or covariance with investment opportunities (i.e., the time-varying hedge component studied in Scruggs (1998) and Rossi and Timmermann (2015), among others). Recent literature on volatility-timing finds that a number of traded factors, and most importantly momentum, have higher returns when volatility is low (Barroso and Santa-Clara (2015) and Muir and Moreira (2017)). In contrast to these findings, which are puzzling in light of conventional wisdom and asset pricing theory, our state variable risk premiums are increasing in conditional variance. We argue that this effect is mostly driven by general macroeconomic uncertainty. In this way, we contribute to a recent literature that analyzes whether risk premiums rise in uncertainty, as implied by long-run risk models. We do not aspire to contribute to the growing literature on measuring uncertainty (see Jurado et al. (2015) and Baker et al. (2016)). Rather we use the conditional variance of consumption as a simple proxy. The asset pricing implications of this type of uncertainty are well understood theoretically and a growing body of empirical work shows that exposure to this risk is priced unconditionally (see, e.g., Boguth and Kuehn (2013), Bansal et al. (2014), and Campbell et al. (2017)).

Naturally, our paper also contributes to literature that tests whether unconditional risk premiums for a range of traded factors as well as state variables are consistent with the ICAPM. Fama and French (1996) advocate an ICAPM interpretation for the size and book-to-market factors of Fama and French (1993). Maio and Santa-Clara (2012) consider various popular multifactor asset pricing specifications and find that most models are not consistent with an ICAPM interpretation.[5] Cooper and Maio (2015) focus on models with multiple traded factors and can only establish best convergence with the ICAPM, as none of the models studied satisfies all the restrictions imposed by the ICAPM. Barbalau, Robotti, and Shanken (2015) develop a formal test for sign consistency and find that with a few exceptions, they cannot reject the null of consistency of the considered models with the ICAPM restrictions. These papers focus on small sets of portfolios as test assets and primarily analyze whether the factors predict aggregate stock market returns in the time series. The approach in Boons (2016) is different, as he finds that risk premiums in a large cross section of individual stocks are consistent with how each state variable predicts macroeconomic activity in the time series. As argued in Cochrane (2005, Ch. 9), macroeconomic variables potentially better proxy for the return on aggregate wealth than stock returns.

Given this conflicting unconditional evidence, it is important to ascertain that conditional risk premiums are consistent with the ICAPM. Chen et al. (1986) and Ferson and Harvey (1991) represent early contributions in this spirit, but these authors focus on small sets of

[5]Lutzenberger (2014) reaches similar conclusions for the European stock market.

test assets and do not decompose the time-variation in an ICAPM risk premium into its relevant economic components. We finally show that the time-variation in state variable risk premiums is not subsumed by the benchmark traded factors of Fama and French (1993) and Fama and French (2015). This suggests that our approach of looking directly at the state variables that theory suggests should capture a time-varying risk premium is attractive. A lot of interesting variation in expected returns is overlooked when analyzing traded factors that proxy only indirectly for state variable risk.

The rest of this paper is organized as follows. Section 1 provides the theoretical motivation for the conditional asset pricing implications of state variable risk. Section 2 tests whether state variables predict consumption growth in a time-varying way. Section 3 presents the method to estimate state variable risk premiums. Section 4 analyzes whether these risk premiums are time-varying. Section 5 analyzes the link with macroeconomic uncertainty. Section 6 concludes. Additional details regarding the bootstrapped standard errors and the simulation analysis are in Appendix A and Appendix B, respectively.

1 Motivation

In this section, we motivate time-variation in state variable risk premiums through a conditional sign restriction that links predictability in the time series of consumption growth to risk premiums in the cross section of stock returns. Consider the conditional asset pricing model $E_t(m_{t+1}r_{t+1,i}) = 0$, where $r_{t+1,i} = R_{t+1,i} - R_{t+1,f}$ is the excess return of asset i; m_{t+1} is a non-negative SDF implied by the absence of arbitrage opportunities; and the expectation is conditional on the investor's information set at time t. Although in most equilibrium models the SDF is a non-linear function of factors and the model's parameters, we follow Cochrane (2005, Ch. 9) and approximate the SDF as a linear function of factors with time-varying coefficients:

$$m_{t+1} = 1 - \boldsymbol{b}_t' \boldsymbol{f}_{t+1}^*. \tag{1}$$

Throughout the paper, we present vectors and matrices in bold. Without loss of generality, we normalize the SDF using zero-mean conditional factors $\boldsymbol{f}^*_{t+1} = \boldsymbol{f}_{t+1} - E_t(\boldsymbol{f}_{t+1})$, such that $E_t(m_{t+1}) = 1$. We assume the vector of factors contains the unexpected excess return of the market portfolio and innovations in a set of K state variables ($\varepsilon_{t+1,z_k} = z_{t+1,k} - E_t(z_{t+1,k})$ for $k = 1, \ldots, K$), such that $\boldsymbol{f}^*_{t+1} = (r^*_{t+1,m}, \boldsymbol{\varepsilon}'_{t+1,z})'$ and $\boldsymbol{b}_t = (b_{t,m}, \boldsymbol{b}'_{t,z})'$.

First and foremost, we test in this paper the hypothesis that $b_{t,z_k} > 0$ when $z_{t,k}$ forecasts the macroeconomy with a positive sign, and vice versa. This hypothesis is motivated by the intuition that investors wish to hedge their risk exposure to state variables that contain systematic economic news. Throughout the paper, and in the following discussion, our main proxy for systematic risk is consumption growth, but the rationale and our empirical evidence extends to alternative measures of macroeconomic activity. Given that marginal utility is lower in good states of the world, assets that pay off when $z_{t,k}$ signals high future consumption growth are not attractive to hedge and have low prices (or high expected returns). This prediction is common to many equilibrium models and relies only on time-separable utility and the assumption that marginal utility is decreasing in consumption.

In our empirical analysis, we measure the coefficients in $\boldsymbol{b}_t$ as the conditional relation between a state variable and future consumption growth estimated over a backward-looking rolling window. This approach allows us to test the conditional asset pricing implications from the model in Eq. (1) using only information that is known at time t. An alternative approach is to assume that the SDF-coefficients, $\boldsymbol{b}_t$, vary over time as a linear function of instruments (see Cochrane (2005, Ch. 8)). However, this approach is designed to estimate the unconditional asset pricing implications from the conditional model in Eq. (1). Another advantage of our approach is that it does not require ex ante knowledge of the instruments that guide the time-variation.

Eq. (1) implies a beta asset pricing model identical to the ICAPM of Merton (1973, Eq. 32):

$$E_t\left(r_{t+1,i}\right) = \lambda_{t,m}\beta_{t,i,m} + \boldsymbol{\lambda}'_{t,z}\boldsymbol{\delta}_{t,i,z}, \tag{2}$$

where $E_t\left(r_{t+1,i}\right)$ is the expected excess return of asset i and $\boldsymbol{\lambda}_t = (\lambda_{t,m}, \boldsymbol{\lambda}'_{t,z})' = Var_t(\boldsymbol{f}^*_{t+1})\boldsymbol{b}_t$ is the $K+1$-vector of risk premiums for exposure to the market ($\beta_{t,i,m}$) and the state variables ($\boldsymbol{\delta}_{t,i,z}$). We analyze the conditional implications from this model by conducting cross-sectional regressions of stock returns on exposures in each period $t+1$. These exposures $\beta_{t,i,m}$ and $\boldsymbol{\delta}_{t,i,z}$ are the time t slope coefficients from the return-generating process

$$r_{i,s+1} = \alpha_{i,t} + \beta_{t,i,m} r_{m,s+1} + \boldsymbol{\delta}'_{t,i,z}\boldsymbol{\varepsilon}_{z,s+1} + \nu_{i,s+1}, \tag{3}$$

where the time index s runs up to $t-1$ to indicate that we use only historical information. As shown in Fama (1976), these cross-sectional risk premium estimates can be interpreted as the return of a zero-investment portfolio with a beta of one to the factor of interest and a beta of zero to all remaining factors in the model. These returns thus recover the risk premiums $\boldsymbol{\lambda}_t$ and estimate the conditional expected excess return investors require to invest in a portfolio with a unit conditional factor beta.

There is already a long literature analyzing time-variation in the market risk premium, $\lambda_{t,m}$, in cross-sectional tests (see, e.g., Jagannathan and Wang (1996), Lettau and Ludvigson (2001), and Lewellen and Nagel (2006)) as well as due to the hedge component of the ICAPM (see, e.g., Scruggs (1998) and Rossi and Timmermann (2015)). Therefore, our main interest is in analyzing the time-variation in the state variable risk premiums, $\boldsymbol{\lambda}_{t,z}$. To ascertain that our results are not driven by a time-varying market risk premium, we orthogonalize the innovations in the state variables from the market in our empirical analysis.

When the innovations in the state variables are also conditionally orthogonal to each other, the state variable risk premiums in $\boldsymbol{\lambda}_{t,z}$ are identical in sign to the respective elements of $\boldsymbol{b}_{t,z}$. Empirically, the off-diagonal elements of $Var_t(\boldsymbol{\epsilon}_{t+1,z})$ are small for the set of state variables we consider. Thus, the model implies a positive conditional risk premium, λ_{t,z_k}, whenever a state variable k currently forecasts consumption growth with a positive sign ($b_{t,z_k} > 0$). Said differently, if the predictive relation between the state variable and consumption growth varies over time, the risk premium must vary over time in a consistent manner. The model further implies that the variation in λ_{t,z_k} due to variation in b_{t,z_k} is magnified when the conditional variance of the state variable, $Var_t(\epsilon_{t+1,z_k})$, is high. As conditional variance increases, shocks to $z_{t,k}$ are more likely to be large in absolute magnitude. Consequently, the premium investors require for a unit beta exposure will also increase.[6] As concluded in Ferson and Harvey (1991), the price of risk per unit of beta with respect to a state variable consists of two components: risk aversion for the state variable, given by b_{t,z_k}, and the conditional variance of the state variable, $Var_t(\epsilon_{t+1,z_k})$.

These testable implications are similar in spirit to two strands of the literature. First, a recent literature analyzes how inflation risk premiums in stock and bond markets vary with the time-varying relation between inflation and consumption growth, or the nominal-real covariance (Bekaert and Wang (2010), Campbell et al. (2014), David and Veronesi (2014), Kang and Pflueger (2015), and Boons et al. (2016)). Second, in ICAPM equilibrium, exposure to state variables that forecast the distribution of future aggregate stock market returns are priced in addition to market beta. The unconditional testable implication is that $b_{z_k} > 0$ when z_k forecasts high returns. As argued by Roll (1977), the aggregate stock market return may be a poor proxy for the return on aggregate wealth, which is the opportunity set of interest to the representative investor. As returns on aggregate wealth are not directly observable, we follow the advice in Cochrane (2005, Ch. 9) and seek instead "recession state variables."

In short, we argue that market beta does not capture all risk that investors care about, as in the ICAPM. Rather, a multifactor model applies, because investors desire to hedge additional risk. In our case, this risk follows from shocks to consumption growth, or macroeconomic growth more generally. We study the model's conditional implications, because the

[6] Although the term $Var_t(\epsilon_{t+1,z_k})$ can be divided out in Eq. (2), it is important empirically. First, it allows us to focus on the risk premium for a fixed unit beta of exposure, instead of covariance, as is common in the literature. Second, it allows for the analysis of commonalities in the variance of the state variables and macro-uncertainty, which is a key determinant of risk premiums in a number of recent asset pricing models.

unconditional implications from ICAPM-type models have been studied recently with mixed and inconclusive results (see, e.g., Maio and Santa-Clara (2012), Cooper and Maio (2015), Barbalau, Robotti, and Shanken (2015), and Boons (2016)). Moreover, these unconditional studies are subject to the critique that unconditional implications do not directly follow from the ICAPM of Merton (1973), which is a conditional model (Hansen and Richard (1987) and Cochrane (2005, Ch. 8)).

2 The time-varying relation between state variables and consumption growth

In this section, we analyze if the relation between state variables and consumption growth is time-varying in an economically important way. As discussed in the previous section, the ICAPM implies that if such time-variation is present, risk premiums for exposure to state variables should vary over time in a consistent manner.

2.1 Data

To start, we measure monthly nominal consumption growth (c_t) using the seasonally-adjusted aggregate nominal consumption expenditures on nondurables and services from the National Income and Product Accounts (NIPA) Table 2.3.5. Population numbers come from NIPA Table 2.1 and price deflator series from NIPA Table 2.3.4, which we use to construct the time series of per capita real consumption growth.

We consider four empirical multi-factor specifications that are popular because the state variables included in each model, denoted $\boldsymbol{z}_t = (z_{t,1}, ..., z_{t,K})'$, have shown strong explanatory power for the cross section and time series. Our first three-factor model contains the default spread (DS) between the yield of long-term corporate BAA- and AAA-rated bonds, the dividend yield (DY) of the CRSP value-weighted stock market index (the ratio of dividends in the previous 12 months to the current level of the index), and the term spread (TS) between the yield of the ten- and one-year Treasury bond. The yield data are taken from

FRED. Second, we follow Petkova (2006) and substitute the risk-free rate for TS. Third, we use the model of Campbell and Vuolteenaho (2004) that includes the price-to-earnings ratio (PE), the value spread (VS), and TS.[7] Fourth, we use the model of Koijen et al. (2017) that includes the bond market factor (CP) of Cochrane and Piazzesi (2005) and a term structure level factor (LVL).[8] Consistent with most empirical studies of the cross section, our sample starts in the second quarter of 1962 and ends in the fourth quarter of 2014.

2.2 Unconditional versus conditional consumption growth predictability

Using these four sets of state variables, we compare the performance of unconditional versus conditional models for predicting consumption growth. First, we run the unconditional predictive regression of log consumption growth (cumulated over horizons $H = 1, 4$) on the state variables from a particular model:

$$c_{t+1:t+H} = b_0 + \boldsymbol{b_z}'\boldsymbol{z}_t + e^u_{t+1:t+H}. \tag{4}$$

Table I presents the results. We see that some state variables predict consumption growth with significant coefficients, such as DY and DS in Model 1. At the four-quarter horizon, the adjusted R^2 ranges from -0.01 in Model 2 to 0.10 in Model 1. Thus, the fraction of the variation in consumption growth explained by the state variables is modest even at the annual frequency.

A possible explanation for this result is that an unconditional regression does not account for time-varying predictability. To see why this is important, we present in Figure 1 rolling ten-year coefficients from a regression of consumption growth on the lagged state variables in each model. For the sake of exposition, we focus on the rolling coefficients for DY, DS,

[7] PE is the cyclically adjusted price-earnings ratio, that is the log ratio of the price of the Standard & Poor's 500 index to the moving-average of its earnings over the previous ten years. VS is calculated from six portfolios sorted on size and book-to-market.

[8] CP is the fitted value from a regression of excess bond returns on forward rates. LVL is the first principal component of the one- through five-year Fama and Bliss forward rates.

and TS in Model 1, RF in Model 2, PE and VS in Model 3, and CP and LVL in Model 4. In short, we see considerable variation in the predictive coefficients for the state variables in each of the four models.

To analyze whether this time-variation in the predictive coefficients is economically important, and not spurious, we follow the two-stage approach introduced in Boons et al. (2016):

$$c_{t+1:t+H} = d_0^H + d_1^H(a_{t,z}^H + \boldsymbol{b}_{t,z}^{H\prime}\boldsymbol{z}_t) + e_{t+1:t+H}^c, \text{ where} \tag{5}$$

$$c_{s+1:s+H} = a_{t,z}^H + \boldsymbol{b}_{t,z}^{H\prime}\boldsymbol{z}_s + e_{s+1:s+H}. \tag{6}$$

In the first stage (Eq. (6)), we regress consumption growth on the lagged state variables over a backward-looking 10-year rolling window using all data available up to quarter t (as in Figure 1).[9] Hence, the window s runs from $t-40-H+1$ to $t-H$, for $H=1,4$. In the second stage (Eq. (5)), we use the estimated coefficients and the state variables observed at time t, i.e., $a_{t,z}^H + \boldsymbol{b}_{t,z}^{H\prime}\boldsymbol{z}_t$, to predict consumption growth from month $t+1$ to $t+H$. This setup ensures that we use no forward-looking information when we predict consumption growth in the second stage.

If this structure correctly models the conditional expectation of consumption growth, we should find that $d_0^H = 0$ and $d_1^H = 1$. However, the estimated $b_{t,z}^H$ contain noise and are subject to misspecification and omitted variables bias. Consequently, one expects to find $d_0^H > 0$ and $d_1^H < 1$ empirically. In other words, it is unlikely that the simple rolling regression fully captures the information set used by investors to form their conditional expectations. However, an interesting question is whether controlling for the rolling relation between the state variables and consumption growth improves the forecast of consumption growth relative to an unconditional model. To this end, we calculate an R^2-comparison

[9] We extend the sample back by an additional five years by estimating the same regression over an expanding rolling window from five to ten years from 1967 to 1972. We fix the length of the rolling window to five and fifteen years in a robustness check.

as: $1 - var(e^c_{t+1:t+H})/var(e^u_{t+1:t+H})$. Finally, we ascertain that the predictive performance does not come from the time-varying intercept, $a^H_{t,z}$, alone, which measures lagged average consumption growth. To this end, we present an R^2 comparison between our model and a specification of Eqs. (5) and (6) that leaves out the state variables. Table II presents the results. To test significance, we report bootstrapped t-statistics that address the concern that our estimates are biased due to errors-in-variables (EIV). Appendix A contains a detailed description of the bootstrap experiment. We find that bootstrapped standard errors lead to conservative inference relative to asymptotic Newey-West standard errors.

At the one-quarter horizon ($H = 1$), d^H_1 is significantly larger than zero for each model at about 0.65. Consistent with this large and significant coefficient, we find that the simple R^2 is much larger than in the unconditional models, ranging from 0.16 to 0.29. This finding implies that consumption growth is predictable and a large chunk of this predictability comes from time-varying coefficients on the state variables.[10] Indeed, the (conditional versus unconditional) model comparison R^2 ranges from 0.10 to 0.28. The R^2 values for the comparison to the conditional model that leaves out the state variables are similarly large. The results at the annual horizon ($H = 4$) are qualitatively similar, though slightly weaker, which suggests that the two-stage predictive approach worsens as the horizon increases. At both horizons, however, we can conclude that the state variables predict consumption growth in an economically important, time-varying way. In the following section, we ask whether this time-variation is mirrored in the risk premium investors desire for exposure to these state variables.

Going forward, we assume the estimates $\widehat{\boldsymbol{b}^H_t} = (\widehat{b^H_{t,z_1}}, ..., \widehat{b^H_{t,z_K}})$ are proportional to the vector of time-varying SDF-coefficients $\boldsymbol{b}_{t,z} = (b_{t,z_1}, ..., b_{t,z_k})$ from Eq. (1). Consequently, we have a simple, backward-looking, empirical proxy that should guide time-variation in the state variable risk premiums in an ICAPM (see Section 1). To ascertain that our results are

[10]This conclusion is robust to including lags of consumption growth. These results are available upon request.

not sensitive to our choice of proxy and to respond to the advice in Cochrane (2005, Ch. 9) to search for "recession state variables," we substitute alternative measures of macroeconomic activity for consumption growth in Eqs. (5) and (6) in a robustness check.[11]

3 Measuring state variable risk premiums

Following recent literature we estimate risk premiums by running stock-level cross-sectional regressions of quarterly returns on historical exposures to innovations in state variables. Among others, Litzenberger and Ramaswamy (1979) and Ang et al. (2016) contend that stock-level tests are more efficient than portfolio-level tests. While tests with individual stocks are plausibly more subject to an error-in-variables problem (due to noise in estimated exposures), the much wider cross section of true exposures should more than offset this shortcoming. For the sake of robustness, we also present state variable risk premiums estimated by sorting stocks into portfolios. In the following, we present a brief overview of our method, which is identical to Boons (2016). We refer the interested reader to his paper for more detail and evidence on the robustness of this approach.

3.1 Innovations as risk factors

Following Campbell (1996), state variables are assumed to follow a first-order VAR. The VAR(1) uses only historical data up to period t so that the betas are fully conditional. Therefore, for each collection comprising the stock market portfolio and a given set of state variables, $\boldsymbol{f}_t = (r_{t,m}, \boldsymbol{z}_t')'$, we estimate $\boldsymbol{f}_s = \boldsymbol{A}_{t,0} + \boldsymbol{A}_{t,1}\boldsymbol{f}_{s-1} + \boldsymbol{\varepsilon}_s$ for $s = 1, \ldots, t$. The residuals $\boldsymbol{\varepsilon}_s$ are orthogonal to the market return, $r_{t,m}$, and we scale them to have the same variance as $r_{t,m}$. From this point forward, $\boldsymbol{\varepsilon}_s$ denotes the transformed innovations that are used as risk factors in the asset pricing model in period t. The correlations between these risk factors are generally weak (< 0.20). Thus, there is no need to orthogonalize the state variables from

[11] For this exercise, we use seasonally-adjusted industrial production growth and real GDP growth, downloaded from the *FRED*® database of the St. Louis FED.

each other, and our tests will provide direct evidence on the conditional consistency of the sign of the cross-sectional risk premium λ_{t,z_k} and that of the b_{t,z_k} coefficients in the SDF. Recall also that the effect of time-variation in b_{t,z_k} is magnified by the conditional variance of the state variables. To consistently measure this conditional variance, denoted σ^2_{t,z_k}, we estimate a GARCH(1,1)-model for the VAR-residuals. Our conclusions are robust to using instead the squared innovations.[12]

3.2 Conditional exposures

Following standard practice, we use all common stocks listed on NYSE, Amex, and Nasdaq, excluding financials and firms with negative book equity. We estimate historical loadings on state variable risk in each period t for all stocks for which we have at least four out of the last five years of returns available. Using only past data that are available each moment in time ensures the analysis is out-of-sample. Elton et al. (1978) and Cosemans et al. (2016), among others, show that loadings on non-traded factors are often small and hard to estimate. Therefore, we use a weighted least squares (WLS) regression over all observations $s = 1, \ldots, t$ to estimate a stock's market beta, $\beta_{t,i,m}$, and exposure to the state variables, $\boldsymbol{\delta}_{t,i,z}$ (see Eqs. (7) and (8) in Boons (2016)). We use an expanding window of returns with an exponential decay in the weights. The expanding window ensures that as much information as possible is used, and an exponential decay in the weights ensures the timeliness of the estimated betas, which is important as loadings of individual firms can vary over time. We shrink these betas to the cross-sectional mean following Vasicek (1973).[13]

[12]These results are available upon request.

[13]An alternative approach is to estimate stock betas by dividing the conditional covariance of returns and innovations in the state variables by the conditional variance of the state variables from a multivariate GARCH-model (to be consistent with our GARCH-measure of conditional variance, σ^2_{t,z_k}). However, applying this approach in an out-of-sample context is next-to-impossible, because such models are hard to estimate using only a small sample of historical data and for many stocks we do not have a long history of returns.

3.3 Cross-sectional regressions

Our main interest in this paper is in analyzing conditional state variable risk premiums $\boldsymbol{\lambda}_{t,z}$, which we estimate in each period t by running a Fama and MacBeth (1973) cross-sectional regression of quarterly returns on lagged historical market beta, $\beta_{t,i,m}$, and loadings on innovations of the state variables of a particular model, $\boldsymbol{\delta}_{t,i,z}$:

$$r_{t+1,i} = \lambda_{t,0} + \lambda_{t,m}\beta_{t,i,m} + \boldsymbol{\lambda}'_{t,z}\boldsymbol{\delta}_{t,i,z} + \upsilon_{t+1,i},\ i = 1, \ldots, N_t. \tag{7}$$

We estimate the state variable risk premiums for the four multi-factor models introduced in Section 2.1. The estimation of the exposures implies an initial burn-in period of five years. As a result, the sample for our asset pricing tests consists of 191 quarters from the second quarter of 1967 to the fourth quarter of 2014.

Panel A of Table III presents summary statistics for the eight risk premiums of interest, i.e., $\lambda_{t,DY}$, $\lambda_{t,DS}$, and $\lambda_{t,TS}$ from Model 1, $\lambda_{t,RF}$ from Model 2, $\lambda_{t,PE}$ and $\lambda_{t,VS}$ from Model 3, and $\lambda_{t,CP}$ and $\lambda_{t,LVL}$ from Model 4.[14] Recall that these state variable risk premiums represent the return of a mimicking portfolio with unit exposure to the state variable of interest in quarter t (and an exposure of zero to other factors in the model). The unconditional risk premiums for DS, TS, RF, VS, and CP are economically large and significant with absolute Sharpe ratios ranging from 0.27 to 0.48. For comparison, the Sharpe ratio of the CRSP market value weighted portfolio over this sample equals 0.35.

The cross-sectional regression in Eq. (7) is subject to EIV. The independent variables are unobserved exposures for each stock pre-estimated in a time-series regression. This time-series regression, in turn, regresses actual returns on VAR-innovations of state variables that are themselves pre-estimated. EIV likely biases the estimates of risk premiums toward

[14]We include each different state variable only once in our analysis of risk premiums to not bias our results in favor of state variables that are included in more than one benchmark model (e.g., DY and DS are present in Models 1 and 2).

zero but it is important to correct the standard errors when we analyze the variation of the risk premiums over time. To do so, we estimate standard errors in all our conditional asset pricing tests using a bootstrap exercise in which we perform 500 replications of the coefficient estimates using artificial block-resampled data (see Appendix A for more detail). In addition, we sort stocks into equal-weighted or value-weighted quintile portfolios using their conditional exposures to the state variables. We then use the return of the high-minus-low quintile portfolio as an alternative measure of the state variable risk premium. With this alternative estimate some stocks will surely be assigned to the wrong portfolio resulting in conservative inference under the null of cross-sectional pricing ability (Boguth and Kuehn (2013)).

4 Time-variation in state variable risk premiums

In this section, we test the conditional asset pricing implications of state variable risk. First and foremost, the ICAPM (see Section 1) implies that the lagged relation between a state variable and consumption growth predicts the risk premium for that state variable with a positive sign. To get some intuition, we start by regressing each individual state variable risk premium (compounded over a horizon of four quarters) on a constant and a dummy variable indicating whether the lagged predictive coefficient for consumption growth, b^4_{t,z_k}, is above the sample median for that state variable:

$$\lambda_{z_k,t+1:t+4} = h_0 + h_1 I_{b^4_{t,z_k}>\text{med}} + e_{z_k,t+1:t+4}. \tag{8}$$

This dummy-specification is attractive as it averages out the measurement error in the estimated b^4_{t,z_k}'s and controls for differences in the unconditional relation between a given state variable and consumption growth. In short, we see that the estimate h_1 is positive for each of the eight state variables under scrutiny. Averaged across state variables, the coefficient estimate is economically large and suggests that risk premiums are increasing in the time series by 4.84% (annualized) when the relation between the state variable and consumption growth is relatively strong. The coefficients are imprecisely estimated, however, as only half of them are individually significant at the 10-% level.

To gain power, we turn to pooled predictive regressions of the form:

$$\lambda_{t+1:t+H,z_k} = g_0 + g_1 b^H_{t,z_k} + e_{t+1:t+H,z_k}, \tag{9}$$

that is, we jointly regress the risk premiums for the state variables on the lagged predictive coefficient between that state variable and consumption growth. Besides the econometric advantage of increasing power, the economic motivation for this pooled specification is that in our model the time-variation in risk premiums is driven by the strength of the relation between a state variable and consumption growth at a given point in time. This hypothesis is not specific to a given state variable nor to a collection of state variables in a particular multi-factor model. Therefore, we run this pooled regression for the complete set of eight state variables. To take care of the noise in the estimated b^H_{t,z_k}'s, we start by substituting in this regression dummy variables either indicating whether a predictive coefficient is positive, $I_{b^H_{t,z_k}>0}$, or whether it is above the sample median (as above). We present t-statistics using both asymptotic standard errors following Driscoll and Kraay (1998), which are robust to very general forms of cross-sectional as well as temporal dependence, and block-bootstrapped standard errors from 500 replications of our estimation procedure, which are typically more conservative (see Appendix A).

Table IV presents the results. In the first column, we run a pooled regression of risk premiums on $I_{b^H_{t,z_k}>0}$ at horizons $H = 1, 4$. At both horizons, the coefficient estimate, g_1, is positive, economically large and statistically significant using both asymptotic and bootstrapped inference. The estimates imply an average increase in a state variable risk premium of 4.90% ($t_{boot} = 2.18$) at $H = 1$ and 6.39% ($t_{boot} = 2.83$) at $H = 4$ when the predictive coefficient between that state variable and consumption growth is positive. This number is relative to a negative risk premium of -2.5%, on average, when the predictive coefficient is negative. Although our focus is on time-variation, this result is also consistent with the idea that a positive (negative) relation between the state variable and consumption growth is rewarded with a positive (negative) state variable risk premium unconditionally.

Therefore, to control for unconditional differences between the predictability of consumption growth from different state variables, we next run a pooled regression of risk premiums on $I_{b^H_{t,z_k}>\text{med}}$. We find that the coefficient estimate, g_1, is similar to the first specification. Qualitatively, this result implies that our estimates capture time-variation in the pool of state variable risk premiums. Quantitatively, the coefficient estimates imply that, on average, when the predictive coefficient for consumption growth for a particular state variable is above its time-series median, the state variable risk premium is larger by an economically large and significant 4.96% ($t_{boot} = 2.88$) for $H = 1$ and 5.17% ($t_{boot} = 2.82$) for $H = 4$.

The next two models run the pooled regression of risk premiums on the raw predictive coefficients, b^H_{t,z_k}, both without and with state variable fixed effects. These two specifications are similar to the first two specifications discussed above, in that the model with state variable fixed effects picks up only time-series variation in the risk premium for a given state variable. We see that our conclusions are robust. The coefficient estimates on b^H_{t,z_k} are (marginally) significant and economically large at both horizons and in both specifications. For instance, the model with fixed effects at $H = 4$ suggests that state variable risk premiums increase over time by about 2.75%, on average, for a standard deviation increase in b^4_{t,z_k}. This number further implies that expected state variable risk premiums move over time about as much as the absolute unconditional average risk premium in the pool of state variables, which is 3.50%.

In the remaining four columns of the table, we run the pooled regression of risk premiums on $I_{b^H_{t,z_k}>\text{med}}$ for each of the four separate models, where Model 1 includes $(\lambda_{t,DY}, \lambda_{t,DS}, \lambda_{t,TS})'$, Model 2 includes $(\lambda_{t,DY}, \lambda_{t,DS}, \lambda_{t,RF})'$, Model 3 includes $(\lambda_{t,TS}, \lambda_{t,PE}, \lambda_{t,VS})'$, and Model 4 includes $(\lambda_{t,CP}, \lambda_{t,LVL})'$. In each model, we see an economically large effect at both horizons, although there is some variation in magnitude and significance. For instance, coefficient estimates range from 2.38% ($t_{boot} = 0.95$) in Model 3 to 7.27% ($t_{boot} = 2.88$) in Model 1 at the annual horizon. Given quantitatively similar effects across the four models, it is only natural that our joint estimates come out as more significant statistically.

4.1 Simulation and robustness checks

Table V presents the results from two Monte Carlo simulation experiments. Appendix B provides a detailed description of each simulation. The first simulation assesses the size of our tests. We ask how often we reject the null hypothesis of no time-variation in an unconditional ICAPM world. In this world, the state variables capture a constant risk premium because each predicts consumption growth with a coefficient that is constant over time. The second simulation assesses the power of our tests. We ask how often we reject the null hypothesis of no time-variation in a conditional ICAPM world. In this world, the state variable risk premiums vary over time, because each state variable predicts consumption growth with a time-varying coefficient. For both simulations, we report the empirical distribution of (i) the R^2-comparison between unconditional and conditional consumption growth predictability (as reported in Table II), (ii) the unconditional risk premiums (as reported in Table III), and (iii) the slope coefficients from the pooled predictive regressions of Eq. (9) (as reported in Table IV). To control for differences in the unconditional state variable risk premiums, we focus on the coefficient estimates from two specifications: the dummy specification with $I_{b^1_{t,z_k} > \text{med}}$ (Spec 2) and the continuous specification with b_{t,z_k} and state variable fixed effects (Spec 4).

The median unconditional average risk premiums in both simulations are the same, but risk premiums are more likely to be significantly positive in the unconditional than in the conditional ICAPM-world (78% versus 61%). In addition, in the unconditional ICAPM world, over 99% of the R^2-comparisons are negative, consistent with the idea that any estimated time-variation in the relation between a state variable and future consumption growth is noise. Hence, it is unsurprising that the estimated slope coefficients are centered around zero and it is in less than 6% of the simulations that the predictive coefficient is positive and significant in either of the two specifications. In contrast, in the conditional ICAPM-world, 60% of the R^2-comparisons are positive, and more than 97% (77%) of the slope coefficients are (significantly) positive. Also, the median estimates are similar to what we find in the data. We conclude that our method does not over-reject and has power to distinguish between the unconditional and conditional ICAPM world.

Next we present in Table VI robustness checks for the method focusing on the pooled regression of Equation (9) with $I_{b^H_{t,z_k}>\text{med}}$, which is a direct and simple test of whether risk premiums vary over time at the state variable level. Panel A analyzes the robustness of our choice of proxy for the time-varying coefficients of the SDF (see Eq. (1)). Before, we measured the conditional relation between a state variable and the consumption-investment opportunity set using its rolling predictive coefficient for consumption growth. Now, we substitute three measures of general macroeconomic activity for consumption growth to estimate b^H_{t,z_k}: GDP growth, industrial production growth, and the first principal component of the three macro-series.[15] In short, we see a positive, economically large and significant coefficient for all three alternatives, indicating that state variable risk premiums increase over time by about 4% when the relation with general macroeconomic activity for that state variable is historically strong. Our results are thus consistent with conditional asset pricing implications from a model where general macroeconomic risk is priced.

Panel B analyzes the robustness of our measures of state variable risk premiums. First, we replace the Fama and MacBeth (1973) state variable risk premium estimates at the stock-level with a high-minus-low quintile risk-sorted portfolio. The risk-sorted portfolios are sorted on exposures ($\boldsymbol{\delta}_{t,i,z}$'s) to innovations in DY, DS, and TS (from Model 1), RF (from Model 2), PE and VS (from Model 3) and CP and LVL (from Model 4). We consider both equal-

[15]The principal component explains about 68% of the common variation in the three series.

weighted (HLEW) and value-weighted (HLVW) portfolios. Furthermore, we estimate the state variable risk premiums at the portfolio-level. The approach is similar to Table IV except that we now use a total of $8 \times 5 = 40$ value-weighted risk-sorted quintile portfolios, instead of individual stocks, to run the quarterly Fama and MacBeth (1973) cross-sectional regressions that estimate the conditional state variable risk premiums, $\boldsymbol{\lambda}_{t,z}$.[16] Using these alternative state variable risk premiums, we run the same pooled regressions as before of returns on the dummy variable indicating whether the conditional predictive coefficient between a state variable and consumption growth is above median. Again, our conclusions are robust, with economically large and typically significant coefficient estimates for g_1 of about 4% for $H = 1$ and 3% for $H = 4$ for each of the three alternative methods.

Panel C analyzes the robustness of our results splitting our sample in two halves at the end of 1990. In this exercise, we define the dummy indicator relative to the median predictive coefficient in each subsample. In this way, we are really asking whether differences in the strength of the relation with consumption growth in a particular subsample capture time-variation in risk premiums. We see that the predictive coefficient is economically large in both sample halves, and significant at the annual horizon at 6.36% ($t = 2.32$) in the first part of the sample and 3.69% ($t = 2.00$) in the second part. We conclude that time-variation in the relation between a state variable and consumption growth robustly predicts state variable risk premiums over time. Panel D shows that our results are not sensitive to the choice of a ten-year rolling window to estimate the conditional relation between consumption growth and the state variables (see Eq. (6)).

Finally, we show in Panel A of Table VII that this time-variation is robust to controlling for contemporaneous exposure of the state variable risk premiums to the benchmark traded factors of the CAPM (Sharpe (1964), Lintner (1965), and Mossin (1966)) and the Fama-French three- and five-factor models (Fama and French (1993, FF3M) and Fama and French

[16]For the sake of comparison, we scale these risk premiums to have the same standard deviation over time as the estimates for individual stocks.

(2015, FF5M)).[17] This finding is interesting given previous literature that argues that these factors (and their underlying characteristics) are associated to unconditionally priced state variable risk (see, e.g., Fama and French (1996), Petkova (2006), and Boons (2016)). We find that the unconditional association between factors and state variables is not enough to capture the time-variation in the price of state variable risk. We conclude that a lot of interesting variation in expected returns might be overlooked when analyzing factors that proxy only indirectly for state variable risk. This is another advantage of our method to focus directly on state variables that theory suggests should capture a time-varying risk premium.

4.2 Conditional variance of the state variables

The ICAPM (see Section 1) implies that the effect of b^H_{t,z_k} on risk premiums strengthens in times of high conditional variance of the state variables. The intuition is that investors want a larger compensation for a unit exposure when there is more risk (in absolute magnitude). To test this implication, we extend the pooled predictive regressions to include an interaction effect:

$$\lambda_{t+1:t+H,z_k} = g_0 + g_1 b^H_{t,z_k} + g_2 \sigma^2_{t,z_k} + g_3 \sigma^2_{t,z_k} \times b^H_{t,z_k} + e_{t+1:t+H,z_k}, \tag{10}$$

where σ^2_{t,z_k} is the GARCH(1,1) conditional variance of a state variable. As before, the dependent variable returns are the risk premiums for the eight state variables of interest. The coefficient g_1 captures the effect on risk premiums of a time-varying relation between the state variable and future consumption. The new coefficients g_2 and g_3 capture the effect of conditional variance on risk premiums. Given that our state variable risk premiums vary in sign (in both the cross section and time series), there is no theoretical guidance on the unconditional effect captured by g_2. The coefficient on the interaction effect, g_3, should be

[17] For this exercise, we orthogonalize the state variable risk premiums using the regression: $\lambda_{t+1,z_k} = \alpha + \boldsymbol{\beta}'\boldsymbol{F}_{t+1} + \epsilon_{t+1}$, where $\boldsymbol{F}_{t+1} = (r_{t+1,m}, r_{t+1,SMB}, r_{t+1,HML})'$ for the FF3M, for instance. From this regression, we use the estimate $\alpha + \epsilon_{t+1}$ for each state variable z_k as the left-hand side risk premium in the pooled regression.

positive, however. If a state variable predicts consumption growth with a positive sign, its risk premium should increase in the conditional variance of the state variable, and vice versa for the negative case.

Given that the three independent variables are substantially correlated, we find that it is hard to interpret the economic and statistical significance of the individual coefficients. To deal with this issue and make reliable inferences on our hypotheses, we estimate the predicted state variable risk premium from the regression in four different cases. These predictions effectively evaluate the joint significance of the coefficients, for which purpose multicollinearity typically has little impact (see, e.g., Judge et al. (1985)). In case one, the predictive relation between the state variable and future consumption growth is weak and the conditional variance of the state variable is high, which case is defined by setting b^H_{t,z_k} equal to one standard deviation below the mean (in the pool) and σ^2_{t,z_k} equal to one standard deviation above the mean (in the pool). In case two, b^H_{t,z_k} is again low, but σ^2_{t,z_k} is at the mean. In case three, b^H_{t,z_k} is one standard deviation above the mean and σ^2_{t,z_k} is at the mean. In case four, both b^H_{t,z_k} and σ^2_{t,z_k} are high. Our model implies that the predicted risk premium should increase monotonically from case one to case four. We also estimate a dummy-specification where we replace b^H_{t,z_k} with the indicator $I_{b^H_{t,z_k}>0}$ and σ^2_{t,z_k} with the indicator $I_{\sigma^2_{t,z_k}>\text{med}}$. In this specification, we should find that $g_2 < 0$ and $g_3 > 0$.[18] Table VIII presents the results for the annual horizon. We present t-statistics for coefficient estimates and predicted risk premiums using both asymptotic Driscoll and Kraay (1998) and block-bootstrapped standard errors (see Appendix A).

Looking at the four cases in Panel B, we see in both specifications that the coefficient estimates jointly imply economically large and statistically significant variation in the predicted risk premiums that is driven by the conditional variance of the state variables. In the continuous specification, the predicted risk premium for a state variable increases mono-

[18]We do not use the indicator $I_{b^H_{t,z_k}>\text{med}}$ here because the interaction effect depends on the sign of b^H_{t,z_k}, not the relative magnitude.

tonically from -4.69% ($t_{boot} = -2.11$) when b^H_{t,z_k} is low and σ^2_{t,z_k} is high (case one) to 4.17% ($t_{boot} = 1.85$) when both b^H_{t,z_k} and σ^2_{t,z_k} are high (case four). In the dummy specification, the increase is even larger from -5.85% ($t_{boot} = -3.17$) to 6.42% ($t_{boot} = 3.48$).[19] The net effect of conditional variance when b^H_{t,z_k} is low (case one versus case two) as well as when b^H_{t,z_k} is high (case three versus case four) is economically large at about 2% in the continuous specification and over 5.5% in the dummy specification. In the dummy specification, this effect is measured by the coefficients g_2 and g_3, which are both significant (using asymptotic as well as bootstrapped standard errors). In the continuous specification, g_3 is positive, as hypothesized, but only significant using asymptotic standard errors.

Panel A of Table IX shows that these results are robust. First, we see quantitatively and qualitatively similar effects when we measure b^H_{t,z_k} using alternative measures of general macroeconomic activity: GDP growth, industrial production growth, and the first principal component of the three macro-series. Second, we have qualitatively and quantitatively similar results when we estimate state variable risk premiums as the High-minus-Low return from a sort into equal- and value-weighted deciles. When we estimate the state variable risk premium using a cross-sectional regression of portfolios, the results are qualitatively similar, but weaker in magnitude. We conclude that time-variation in the conditional variance of the state variables magnifies the time-variation in state variable risk premiums that is due to time-variation in the relation of the state variables with future consumption-investment opportunities.[20]

5 Macroeconomic uncertainty

In this section we analyze whether the results for the conditional variance of the state variables are driven by state variable specific variance or rather associated to general macroe-

[19] In both specifications, the difference between these two extreme cases is about 10% ($t_{boot} > 3$) and measures the total variation in state variable risk premiums.

[20] Panel B of Table VII shows that controlling for contemporaneous exposure to the CAPM and Fama-French factors does not materially affect this conclusion.

conomic uncertainty. If state variables contain important information about future macroeconomic activity, and the conditional variance of the state variables moves in a correlated manner over time, this will translate into time-varying macroeconomic uncertainty. Hence, all state variable risk premiums could rise during times of high economic uncertainty. Similarly, the risk premium for exposure to news about future consumption growth, and thus the market risk premium, increases in uncertainty in the long-run risk model of Bansal and Yaron (2004). We measure macroeconomic uncertainty as the variance of consumption growth conditional on the state variables. Previous literature finds that related proxies of uncertainty are priced unconditionally in various asset classes. Different from these papers, we ask whether uncertainty drives time-variation in state variable risk premiums and to what extent uncertainty subsumes the role of time-varying variance of the state variables.

5.1 Time-varying macroeconomic uncertainty and state variable risk premiums

To start, we ask whether state variable risk premiums vary over time with the conditional variance of consumption. From the two-stage setup for consumption growth predictability in Eqs. (5) and (6), we define two measures of conditional variance by estimating a GARCH(1,1)-model for the innovations:

$$e^1_{t+1:t+H} = c_{t+1:t+H} - \{d^H_0 + d^H_1(a^H_{t-1,z} + \boldsymbol{b}^{H\prime}_{t-1,z}\boldsymbol{z}_t)\} \text{ and} \tag{11}$$

$$e^2_{t+1:t+H} = c_{t+1:t+H} - \{a^H_{t-1,z} + \boldsymbol{b}^{H\prime}_{t-1,z}\boldsymbol{z}_t\}. \tag{12}$$

We denote the GARCH conditional variance estimates $\sigma^{2,*}_{t,z,H}$ and $\sigma^{2,**}_{t,z,H}$. These measures ensure consistency in the way we measure the conditional expectation and variance of consumption growth. To assess the sensitivity of our results, we also construct a third measure of conditional variance by estimating a GARCH(1,1)-model for de-meaned consumption growth, $e^3_{t+1:t+H} = c_{t+1:t+H} - \mu_{t+1:t+H,c}$, which we denote $\sigma^{2,***}_{t,H}$. Note that this third measure does not depend on the collection of state variables in a particular model.

Next, we substitute each of the three uncertainty measures, $\sigma^2_{t,H}$, for the conditional variance of the state variables in the pooled predictive regressions of Eq. (10) and estimate:

$$\lambda_{t+1:t+H,z_k} = g_0 + g_1 b^H_{t,z_k} + g_2\sigma^2_{t,H} + g_3\sigma^2_{t,H} \times b^H_{t,z_k} + e_{t+1:t+H,z_k}. \tag{13}$$

As before, we predict that the coefficient on the interaction effect, g_3, is positive: If a state variable predicts consumption growth with a positive sign, its risk premium should increase with uncertainty, and vice versa for the negative case. To deal with the correlation between the independent variables and facilitate economic and statistical inference, we again (i) provide the predicted risk premium in the four separate cases (Low b^H_{t,z_k}, High $\sigma^2_{t,H}$; Low b^H_{t,z_k}, Low $\sigma^2_{t,H}$; High b^H_{t,z_k}, Low $\sigma^2_{t,H}$; High b^H_{t,z_k}, High $\sigma^2_{t,H}$) and (ii) estimate a dummy specification where we replace b^H_{t,z_k} with the indicator $I_{b^H_{t,z_k}>0}$ and $\sigma_{t,H}$ with the indicator $I_{\sigma^2_{t,H}>\text{med}}$. Table X presents the results for the annual horizon. We present t-statistics for coefficient estimates and predicted risk premiums using both asymptotic Driscoll and Kraay (1998) and block-bootstrapped standard errors (see Appendix A).

Looking at the four cases in Panel B, we see in both specifications that the coefficient estimates jointly imply economically large and statistically significant variation in the predicted risk premium for state variables driven by the conditional variance of consumption growth. In the continuous specification, the predicted risk premium increases monotonically from -4.81% ($t_{boot} = -2.29$) when b^H_{t,z_k} is low and σ^2_{t,z_k} is high (case one) to 4.56% ($t_{boot} = 2.03$) when both b^H_{t,z_k} and σ^2_{t,z_k} are high (case four). In the dummy specification, the increase is similar from -4.20% ($t_{boot} = -2.39$) to 5.00% ($t_{boot} = 3.01$).[21] The net effect of macroeconomic uncertainty when b^H_{t,z_k} is low (case one versus case two) as well as when b^H_{t,z_k} is high (case three versus case four) is economically large at about 2% in the continuous specification and about 3% in the dummy specification. In the dummy specification, this

[21]In both the continuous and dummy specification, the difference between these two extreme cases is about 9% ($t_{boot} > 3$) and measures the total variation in state variable risk premiums.

effect is measured by the coefficients g_2 and g_3, which are both marginally significant (using asymptotic as well as bootstrapped standard errors). In the continuous specification, g_3 is positive, as hypothesized, but only significant using asymptotic standard errors.

In the remainder of the table, we see robust evidence for the alternative measures of macroeconomic uncertainty: $\sigma^{2,**}_{t,z,4}$ and $\sigma^{2,***}_{t,4}$. Panel B of Table IX further shows that these results are robust to using alternative measures of macroeconomic activity (instead of consumption growth) to measure b^{H}_{t,z_k} and $\sigma^{2,*}_{t,z,4}$ as well as when we estimate state variable risk premiums using portfolios (instead of a cross-sectional regression of individual stocks). We conclude that macroeconomic uncertainty, as measured by the variance of consumption growth conditional on the state variables, magnifies time-variation in state variable risk premiums.[22]

5.2 Macroeconomic uncertainty and the conditional variance of state variables

In this subsection, we analyze to what extent the results for the conditional variance of the state variables (in Section 4.2) are driven by macroeconomic uncertainty. We have shown that the two variance-measures affect state variable risk premiums in a similar way, and the measures are, in fact, correlated. Figure 2 plots the conditional variance of consumption growth ($\sigma^{2,***}_{t,H}$) versus the conditional variance of the eight state variables. We aggregate these conditional variances by taking an average across state variables ($\sigma^{2}_{t,z_k,AVG}$) and by estimating the first principal component ($\sigma^{2}_{t,z_k,PC1}$). The correlation between $\sigma^{2,***}_{t,H}$ and these measures of conditional state variable variance is about 0.5. Moreover, in about 60% of the quarters that $\sigma^{2,***}_{t,H}$ is above median, we also find that $\sigma^{2}_{t,z_k,AVG}$ and $\sigma^{2}_{t,z_k,PC1}$ are above median.[23] This evidence suggests that periods with high common variance of the state variables also tend to be periods with high conditional variance of consumption growth.

[22]Panel C of Table VII shows that controlling for contemporaneous exposure to benchmark asset pricing factors does not materially affect this conclusion.

[23]This fraction is similar when we do the comparison using the two alternative measures of the conditional variance of consumption growth, $\sigma^{2,*}_{t,H}$ and $\sigma^{2,**}_{t,H}$.

Hence, a large share of the variation in risk premiums due to conditional variance of the state variables could be driven by macroeconomic uncertainty.

To test this hypothesis, we repeat the regressions of Section 4.2 using measures for the conditional variance of the state variables that control for quarters where macroeconomic uncertainty is high. To do so, we set the conditional variance of a state variable, σ^2_{t,z_k}, equal to its time-series median in quarters when macro-uncertainty, $\sigma^2_{t,z,H}$, is above median. In the dummy specification, we set $I_{\sigma^2_{t,H}>\text{med}}$ to zero in quarters when $\sigma^2_{t,z,H}$ is above median. To proxy for macro-uncertainty, we consider each of the three measures of the conditional variance of consumption growth. Table XI presents the results, where, for the sake of brevity, we focus on the predicted risk premiums in the four cases (from low b^H_{t,z_k} and high σ^2_{t,z_k} (controlling for macroeconomic uncertainty) in case one to high b^H_{t,z_k} and High σ^2_{t,z_k} (controlling for macroeconomic uncertainty) in case four. We present t-statistics for coefficient estimates and predicted risk premiums using both asymptotic Driscoll and Kraay (1998) and block-bootstrapped standard errors (see Appendix A).

For the first measure of macroeconomic uncertainty (columns one and four), we see that the effect of conditional variance of the state variables is reduced by about two-thirds relative to what we saw before. For the continuous specification, the difference in predicted risk premium when variance is high versus low and when b^H_{t,z_k} is low (high) is -2.89-(-2.35)=-0.55 (3.30-2.56=0.74%) when controlling for macroeconomic uncertainty. In Table VIII these same differences were -2.56% and -1.99%, respectively. For the dummy specification, the effect of controlling for macro uncertainty is even more pronounced at differences between high and low variance of -0.95% when b^H_{t,z_k} is low and 1.70% when b^H_{t,z_k} is high. These numbers are relative to absolute differences of over 5.5% in Table VIII. In the remaining columns of the table we see that these conclusions extend to alternative measures for the conditional variance of consumption growth. Moreover, Panel C of Table IX shows that these results are robust to using alternative measures of macroeconomic activity as well as when we estimate state variable risk premiums using portfolios. We conclude that the bulk of the time-variation in state variable risk premiums due to time-variation in variance is not state variable-specific, rather it is general macroeconomic uncertainty that matters most.

6 Conclusion

In this paper, we show that risk premiums for exposure to state variables are substantially time-varying. We argue that this variation is driven by time-variation in the predictive relation between the state variables and future consumption growth and magnified by time-variation in the conditional variance of consumption growth.

More specifically, we show that controlling for time-varying coefficients in the relation between the state variables and consumption growth significantly improves consumption growth predictability relative to an unconditional model. Consistent with this time-variation, we show the risk premium for a state variable is larger in the time series by about 0.35 in Sharpe ratio whenever a state variable predicts consumption growth relatively strongly. This effect is statistically significant, robust, and economically large: state variable risk premiums move over time about as much as the market risk premium. Further, we show that this effect doubles in strength when the conditional variance of a state variable is high. We argue that an important share of the variation in conditional variances is common to all state variables and associated to general macroeconomic uncertainty, measured as the variance of consumption growth conditional on the state variables.

Our evidence is consistent with an SDF that prices systematic economic news and, in particular, an ICAPM. We contribute to recent literature that tests the unconditional implications of ICAPM-type models, as the evidence so far is mixed and inconclusive. Using the broadest possible cross section of individual stocks, a large set of state variables, and a conditional setup, we raise the bar considerably for the data to be consistent with the model. Our evidence showing that time-variation in risk premiums is magnified by the time-variation in the conditional variance of the state variables contributes to a large literature on the time-varying risk-return trade-off. The fact that a large share of the variation in conditional variances translates to uncertainty about future consumption growth adds to the asset pricing implications of general macroeconomic uncertainty.

Our results on the time-varying relation between state variables and consumption growth contribute to recent literature that builds a time-varying price of inflation risk into popular asset pricing models to address the observed time-variation in the relation between inflation and the real economy. The evidence for state variables related to the level and slope of the yield curve is broadly consistent with these models. Having said that, future work is needed to pin down exactly why our state variables are predicting consumption growth in a time-varying way. Natural candidates for factors driving this variation are monetary policy, the business cycle, and the prevalence of supply and demand shocks. In either case, our findings strengthen the case for time-varying, rational risk premiums that drive the predictability of returns.

Appendix A Bootstrap

The block-bootstrap algorithm designed to deal with EIV consists of the following steps. First, in each of 500 replications, $m = 1, \ldots, 500$, we construct pseudo-samples for consumption growth, the state variables, and the state variable risk premiums. We draw with replacement T_m overlapping three-year blocks from:

$$\{c^m_{t+1:t+12}, \boldsymbol{z}^m_{t+1:t+12}, \boldsymbol{\lambda}_{t+1:t+12,z}\}, \ t = s^m_1, s^m_2, \ldots, s^m_{T_m} \tag{A.1}$$

where the time indices, $s^m_1, s^m_2, \ldots, s^m_{T_m}$, are drawn randomly from the original time sequence $1, \ldots, T$. The block size is chosen to preserve the (auto-)correlation structure in the data and to respect the estimation setup in Eqs. (5) and (6) of the paper. Note, we do not re-estimate the state variable risk premiums by re-sampling the CRSP file. Rather, we re-sample directly from the estimated state variable risk premiums. First, this reduces the computing time of the bootstraps significantly. Second, Boons (2016) already performs a number of simulation and bootstrap exercises to analyze the robustness of the state variable risk premiums to EIV (and misspecification).[24] We join these blocks to construct a quarterly time series matching the length of our sample from the second quarter of 1967 to the fourth quarter of 2014.

For each replication, m, we then run the two-stage tests described in Section 2 for the artificial data:

$$c^m_{t+1:t+H} = d^{m,H}_0 + d^{m,H}_1 (a^{m,H}_{t-1} + \boldsymbol{b}^{m,H\prime}_{t-1} \boldsymbol{z}^m_t) + e^{m,c}_{t+1:t+H}, \text{ where} \tag{A.2}$$

$$c^m_{s+1:s+H} = a^{m,H}_{t-1} + \boldsymbol{b}^{m,H\prime}_{t-1} \boldsymbol{z}^m_s + e^m_{s+1:s+H}, \tag{A.3}$$

and save the estimates $d^{m,H}_0$, $d^{m,H}_1$, and $\boldsymbol{b}^{m,H}_{t-1}$, for $H = 1, 4$. The bootstrapped standard errors for Table II are calculated as the standard deviation of $d^{m,H}_0$ and $d^{m,H}_1$ over the 500 bootstrap replications. The bootstrap estimates, $\boldsymbol{b}^{m,H}_{t-1}$, are going to be used to get the bootstrapped standard errors for the remaining tests of the paper. For these tests, we also bootstrap the conditional variance of consumption, denoted $\sigma^{2,m,*}_{t,z,H}$ by estimating a GARCH(1,1)-model for the innovations $e^{m,c}_{t+1:t+H}$. As in the paper, we also construct the two alternative measures of the conditional variance of consumption: $\sigma^{2,m,**}_{t,z,H}$ (a GARCH(1,1) conditional variance for the innovation $c^m_{t+1:t+H} - (a^{m,H}_{t-1} + \boldsymbol{b}^{m,H\prime}_{t-1} \boldsymbol{z}^m_t))$ and $\sigma^{2,m,***}_{t,z,H}$ (a GARCH(1,1) conditional variance for de-meaned consumption growth, $c^m_{t+1:t+H} - \mu_{c^m_{t+1:t+H}}$). Following Section 3.1, we also construct bootstrap estimates for the GARCH(1,1) conditional variance of the state variables, denoted $\sigma^{2,m}_{t,z_k}$.

Using the bootstrapped predictor data, we then run the predictive regression described in Sections 4 and 5 of the paper. As in Table IV, we first regress the artificial state variable risk premiums (compounded over horizons $H = 1, 4$ quarters) on the lagged predictive coefficient

[24] We found consistent evidence in a bootstrap of the CRSP file, although using only a small number of replications.

$b_{t-1}^{m,H}$ (or the dummy variables derived from these coefficients) using:

$$\lambda_{t+1:t+H,z_k}^{m} = g_0^m + g_1^m b_{t,z_k}^{m,H} + e_{t+1:t+H,z_k}^{m}. \tag{A.4}$$

We also regress the artificial state variable risk premiums (compounded over horizons $H = 1, 4$ quarters) on the lagged predictive coefficient $b_{t-1}^{m,H}$ (or the dummy variables derived from these coefficients), a measure of conditional variance, and its interaction using:

$$\lambda_{t+1:t+H,z_k} = g_0^m + g_1^m b_{t,z_k}^{m,H} + g_2^m \sigma_{t,H}^{2,m} + g_3^m \sigma_{t,H}^{2,m} \times b_{t,z_k}^{m,H} + e_{t+1:t+H,z_k}^{m}, \tag{A.5}$$

where $\sigma_{t,H}^{2,m}$ is the conditional variance of a state variable (as in Table VIII), the conditional variance of consumption (as in Table X) or a measure of the conditional variance of a state variable orthogonalized from the conditional variance of consumption growth (as in Table XI). The timing in the different steps of the bootstrap is consistent with the data so that the timing of the left-hand side returns is strictly after the timing of consumption growth and the state variables used to estimate the right-hand side predictive coefficient. We use the standard deviation of the estimates $(g_0^m, g_1^m, g_2^m, g_3^m)$ over the 500 bootstrap replications as the standard error for the predictive regressions in the paper.

Appendix B Simulation analysis

We conduct two simulations. First, we assume that the unconditional ICAPM holds to assess whether our empirical design has good size properties. Second, we assume that the conditional ICAPM holds to analyze the power properties of our tests. Both simulations use Model 1 as the representative model for the time-series properties of the state variables and are calibrated to the quarterly data. Consumption in Model 1 is a function of DY, DS, and TS, and its choice as a reference model is consistent with the univariate simulations in Maio and Santa-Clara (2012) and Boons (2016) who use TS as the representative state variable. Below we sometimes omit the specific reference to Model 1 when we talk about sample values but stress that the simulations are consistent with the sample values in Model 1. Finally, to allow for a clearer interpretation of our results and more transparency, we use only diagonal co-variance matrices for residuals and use for all three state variables the same slope coefficients in the consumption process, the same parameters determining the time-series behavior of these slopes, and the same unconditional means for the state variable risk premiums.

Both simulations are based on 15,000 samples of artificial data for the three state variables stacked together in $\boldsymbol{z}_{t+1}^{n}$, with $t = 1, \ldots, 210$ and $n = 1, \ldots, 15,000$. State variables follow an AR(1) process

$$\boldsymbol{z}_{t+1}^{n} = \boldsymbol{\phi}_{z}' \boldsymbol{z}_{t}^{n} + \boldsymbol{\epsilon}_{z,t+1}^{n}, \tag{B.1}$$

where the elements in $\boldsymbol{\phi}_z$ are set equal to their respective sample values for DY, DS, and TS, and $\boldsymbol{\epsilon}_{z,t+1}^{n}$ has a multivariate normal distribution with zero means, and a co-variance matrix with diagonal elements of $\mathbf{1} - \boldsymbol{\phi}_z \circ \boldsymbol{\phi}_z$, where $\circ$ is the Hadamard product, and zero off-diagonal elements. The drift in Eq. (B.1) is zero to take into consideration that the state variables are standardized in the analysis. The initial realization of each state variable is drawn for each pseudo-sample from a standardized normal distribution.

In the unconditional ICAPM world, log consumption growth $c_{t+1}^{u,n}$ follows

$$c_{t+1}^{u,n} = b_0^u + \boldsymbol{b}_1^{u\prime} z_t^{u,n} + \epsilon_{c,t+1}^{u,n}, \tag{B.2}$$

where b_0^u is set to its sample value, $\epsilon_{c,t+1}^{u,n}$ is sampled from a normal distribution with zero mean and variance set to its sample value, and all elements in $\boldsymbol{b}_1^u$ are set equal to the average absolute slope coefficient in Model 1 (see Table I).

State variable risk premiums $\boldsymbol{\lambda}_{t+1,\boldsymbol{z}}^{u,n}$ are generated from

$$\boldsymbol{\lambda}_{t+1,\boldsymbol{z}}^{u,n} = \boldsymbol{b}_1^u k + \boldsymbol{\epsilon}_{\boldsymbol{\lambda},t+1}^{u,n}, \tag{B.3}$$

where k is a scaling factor set to the average ratio of risk premiums to slope coefficients in the unconditional consumption predictability regression, and $\boldsymbol{\epsilon}^{u,n}_{\boldsymbol{\lambda},t+1}$ is sampled from a multivariate normal distribution with zero means and diagonal co-variance matrix in which the variances are equal to the average sample variance of the state variable risk premiums of DY, DS, and TS.

To simulate data consistent with a conditional ICAPM, we relax the assumption that the intercepts and slopes in the consumption process are fixed, and instead assume that $b^{c,n}_{0,t+1}$ and $\boldsymbol{b}^{c,n}_{1,t+1}$ follow AR(1) processes. These processes are parametrized to AR(1) processes (henceforth, simply referred to as sample values) fitted to the coefficients of the rolling window regressions of consumption growth on DY, DS, and TS in the empirical analysis. In particular, the intercept $b^{c,n}_{0,t+1}$ follows

$$b^{c,n}_{0,t+1} = \phi_{b_0,0} + \phi_{b_0,1} b^{c,n}_{0,t} + \epsilon^{c,n}_{b_0,t+1}, \tag{B.4}$$

where $\phi_{b_0,0}$ and $\phi_{b_0,1}$ equal their sample values obtained by estimating an AR(1) model on the intercepts in the rolling window regressions of consumption on past state variables, $\epsilon^{c,n}_{b_0,t+1}$ is drawn from a normal distribution with zero mean and variance equal to the variance of the residual in the AR(1) model estimated on the rolling window regressions' intercepts, and $b^{c,n}_{0,1}$ is sampled from the unconditional distribution implied by Eq. (B.4). Similarly, the slope coefficients of the consumption process follow

$$\boldsymbol{b}^{c,n}_{1,t+1} = \boldsymbol{\phi}_{\boldsymbol{b}_1,0} + \boldsymbol{\phi}'_{\boldsymbol{b}_1,1} \boldsymbol{b}^{c,n}_{1,t} + \boldsymbol{\epsilon}^{c,n}_{\boldsymbol{b}_1,t+1}, \tag{B.5}$$

where $\boldsymbol{\phi}_{\boldsymbol{b}_1,1}$ is set to 0.97 which is consistent with the empirical values which range from 0.94 to 0.98, the elements in $\boldsymbol{\phi}_{\boldsymbol{b}_1,0}$ are chosen such that the slope coefficients in the conditional simulation, $\boldsymbol{b}^{c,n}_{1,t+1}$, and unconditional simulation, $\boldsymbol{b}^{u}_{1,t+1}$, have the same expected value, and $\boldsymbol{\epsilon}^{c,n}_{\boldsymbol{b}_1,t+1}$ is drawn from a multivariate normal distribution with zero means and diagonal covariance matrix with variances equal to the average of the sample values. The initial values, $\boldsymbol{b}_{1,0}^{c,n}$, are drawn from the unconditional distribution implied by Eq. (B.5).

Log consumption growth in the conditional ICAPM $c_{t+1}^{c,n}$ is then generated as

$$c_{t+1}^{c,n} = \kappa_0 + \kappa_1 \left(b_{0,t+1}^{c,n} + \boldsymbol{b}_{1,t+1}^{c,n'} \boldsymbol{z}_t^{c,n} \right) + \epsilon_{c,t+1}^{c,n}, \tag{B.6}$$

where κ_0 and κ_1 are two scalars which are chosen such that consumption has the empirically observed mean and the same theoretical variance than in the unconditional simulation, and $\epsilon_{c,t+1}^{c,n}$ is sampled from a normal distribution with mean zero and variance equal to the average variance in the rolling window regressions of Model 1.

State variable risk premiums $\boldsymbol{\lambda}_{t+1,\boldsymbol{z}}^{c,n}$ are generated with

$$\boldsymbol{\lambda}_{t+1,\boldsymbol{z}}^{c,n} = \boldsymbol{b}_{1,t+1}^{c,n} \times k + \boldsymbol{\epsilon}_{\boldsymbol{\lambda},t+1}^{c,n}, \tag{B.7}$$

where the scaling factor k is set to the same value as in the unconditional simulation, and $\boldsymbol{\epsilon}_{\boldsymbol{\lambda},t+1}^{c,n}$ is sampled from a multivariate normal distribution with zero means and diagonal co-variance matrix with variances equal to their average sample value in Model 1.

Armed with the two simulated data sets, $\{\boldsymbol{z}_{t+1}^{n},\ \boldsymbol{c}_{t+1}^{u,n},\ \boldsymbol{\lambda}_{t+1}^{u,n}\}$ and $\{\boldsymbol{z}_{t+1}^{n},\ \boldsymbol{c}_{t+1}^{c,n},\ \boldsymbol{\lambda}_{t+1}^{c,n}\}$, we first run the analysis of Table II and report the R^2 comparison of the conditional versus unconditional model in Table V. Second, we run the analysis of Table III and report the unconditional state variable risk premiums. Finally, we combine the risk premium estimates from three separate simulations, reducing thereby the unconditional and conditional data sets to 5000 simulations each, and run the analysis of Table IV. In this way, the total number of risk premiums in the pooled predictive regression is nine, close to the eight we use in the paper. Because specifications one and three of the pooled predictive regression in that table (using, respectively, the dummy $I_{b_{t,z_k}^4 > 0}$ and the raw measure b_{t,z_k}^4 without fixed effects) cannot distinguish between constant state variable risk premiums with different signs and conditionally priced state variables, we only report specifications two and four (using, respectively, the dummy $I_{b^4_{t,z_k}>\text{med}}$ and the raw measure b^4_{t,z_k} with fixed effects) for the simulation.

References

Ang, A., Liu, J., Schwarz, K., 2016. Using stocks or portfolios in tests of factor models. Unpublished working paper. Columbia University, New York.

Baker, S.R., Bloom, N., Davis, S., 2016, Measuring economic policy uncertainty, *Quarterly Journal of Economics* 131, 1593-1636.

Bansal, R., and A. Yaron, 2004, Risks for the long run: A potential resolution of asset pricing puzzles, *Journal of Finance* 59, 1481-1509.

Bansal, R., D. Kiku, I. Shaliastovich, and A. Yaron, 2014, Volatility, the Macro-economy and Asset Prices, *Journal of Finance*, 69, 2471-2511.

Barbalau, A., Robotti, C., and J. Shanken, 2015, Testing Inequality Restrictions in Multi-factor Asset-Pricing Models, Unpublished working paper.

Barroso, P., and Santa-Clara, P., 2015, Momentum has its moments, *Journal of Financial Economics* 116, 111-120.

Bekaert G., and X. Wang, 2010, Inflation Risk and the Inflation Risk Premium, *Economic Policy* 25, 755-806.

Boguth, O., Kuehn, L., 2013. Consumption volatility risk. *Journal of Finance*, 68, 2589-2615.

Boons, M.F., 2016, State variables, macroeconomic activity and the cross section of individual stocks, *Journal of Financial Economics* 119, 489-511.

Boons, M.F., Duarte, F., de Roon, F.A., Szymanowska, M., 2016, Time-Varying Inflation Risk and the Cross-Section of Stock Returns, Unpublished working paper.

Campbell, J.Y., 1996. Understanding risk and return. *Journal of Political Economy* 104, 298-345.

Campbell, J.Y., Cochrane, J., 1999, By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior, *Journal of Political Economy* 107, 205-251.

Campbell, J.Y., Cochrane, J., 2000, Explaining the Poor Performance of Consumption-based Asset Pricing Models, *Journal of Finance* 55, 2863-2878.

Campbell, J.Y., S. Giglio, C. Polk, and R. Turley, 2017, An Inter-temporal CAPM with Stochastic Volatility, *Journal of Financial Economics*, forthcoming.

Campbell, J.Y., C. Pflueger, and L.M. Viceira, 2014, Monetary Policy Drivers of Bond and Equity Risks, Working Paper Harvard Business School.

Campbell, J.Y., Vuolteenaho, T., 2004. Bad beta, good beta. *American Economic Review* 94, 1249-1275.

Chen, N., Roll, R., Ross, S.A., 1986. Economic forces and the stock market. *Journal of Business* 59, 383-403.

Chordia, T., Goyal, A., Shanken, J., 2015. Cross-sectional asset pricing with individual stocks: betas versus characteristics. Unpublished working paper. Emory University, Atlanta, GA.

Cochrane, J.H., 2005. Asset pricing. Princeton, Princeton University Press.

Cochrane, J.H., Piazzesi, M., 2005. Bond risk premia. *American Economic Review* 95,138-60.

Cooper, I., Maio, P., 2015. Equity risk factors and the intertemporal CAPM. Unpublished working paper. Norwegian Business School, Oslo.

Cosemans, M., Frehen, R., Schotman, P.C., Bauer, R., 2016. Estimating security betas using prior information based on firm fundamentals, *Review of Financial Studies* 29, 1072-1112.

Dangl, T., and M. Halling, 2012, Predictive Regressions with Time-Varying Coefficients, *Journal of Financial Economics* 106, 157-181.

David, A., and P. Veronesi, 2014, What Ties Return Volatilities to Price Valuations and Fundamentals?, *Journal of Political Economy* 121, 682-746.

Driscoll, J. C., Kraay, A.C., 1998, Consistent Covariance Matrix Estimation with Spatially Dependent Panel Data, *Review of Economics and Statistics* 80, 549-560.

Duffee, G., 2005, Time Variation in the Covariance between Stock Returns and Consumption Growth, *Journal of Finance* 60, 1673-1711.

Elton, E.J., Gruber, M.J., Urich, T.J., 1978. Are betas best? *Journal of Finance* 33, 1375-1384.

Fama, E.F., 1976. Foundations of finance. Basic Books, New York.

Fama, E.F., 1991. Efficient capital markets. *Journal of Finance* 46, 1575-1617.

Fama, E.F., French, K.R., 1993. Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics* 33, 3-56.

Fama, E.F., French, K.R., 1996, Multifactor Explanations of Asset Pricing Anomalies, *Journal of Finance* 51, 55-84.

Fama, E.F., French, K.R., 2015, A five-factor asset pricing model, *Journal of Financial Economics* 116, 1-22.

Fama, E.F., MacBeth, J.D., 1973. Risk, return, and equilibrium: empirical tests. *Journal of Political Economy* 81, 607-636.

Ferson, W.E., Harvey, C.R., 1991. The variation of economic risk premiums. *Journal of Political Economy* 99, 385-415.

French, K.R., Schwert, W., and Stambaugh, R., 1987, Expected stock returns and volatility, *Journal of Financial Economics* 19, 3-29.

Gagliardini, P., Ossola, E., Scaillet, O., 2012, Time-Varying Risk Premium in Large Cross-Sectional Equity Datasets, *Econometrica* 84, 985-1046.

Ghysels, E., Santa-Clara, P., and Valkanov, R., 2005, There is a risk-return trade-off after all, *Journal of Financial Economics* 76, 509-548.

Judge, G.G., Griffiths, W.E., Hill, R.C., Lutkepohl, H., Lee, T., 1985, The Theory and Practice of Econometrics, Wiley, New Jersey.

Hansen, L.P., Richard, S.F., 1987, The Role of Conditioning Information in Deducing Testable Restrictions Implied by Dynamic Asset Pricing Models, *Econometrica* 55, 587-613.

Harvey, C.R., 1988, The real term structure and consumption growth, *Journal of Financial Economics* 22, 305-333.

Henkel, S. J., Martin, J. S. and Nardari, F., 2011, Time-Varying Short-Horizon Return Predictability, *Journal of Financial Economics* 99, 560-580.

Jagannathan, R., Z. Wang, 1996, The Conditional CAPM and the Cross-Section of Expected Returns, *Journal of Finance* 51, 3-53.

Jurado, K., Ludvigson, S.C., and Ng, S., 2015, Measuring uncertainty, *American Economic Review* 105, 1177-1216.

Kang, J., and C. Pflueger, 2015, Inflation Risk in Corporate Bonds, *Journal of Finance* 70, 115-162.

Koijen, R., Lustig, H., Van Nieuwerburgh, S., 2017. The cross section and time series of stock and bond returns, *Journal of Monetary Economics* 88, 50-69.

Lettau, M., and S. Ludvigson, 2001, Resurrecting the (C)CAPM: A Cross-Sectional Test When Risk Premia Are Time-Varying, *Journal of Political Economy* 109, 1238-1287.

Lewellen, J.S., Nagel, S., 2006. The conditional CAPM does not explain asset-pricing anomalies. *Journal of Financial Economics* 82, 289-314.

Lewellen, J., Nagel, S., Shanken, J., 2010. A skeptical appraisal of asset pricing tests. *Journal of Financial Economics* 96, 175-194.

Lintner, J., 1965, The Valuation of Risk Assets and the Selection of Risky Investments in Stock Portfolios and Capital Budgets, *Review of Economics and Statistics* 47, 13-37.

Litzenberger, R.H., Ramaswamy, K., 1979. The effect of personal taxes and dividends on capital asset prices. *Journal of Financial Economics* 7, 163-196.

Lutzenberger, F. T., 2014, Multifactor models and their consistency with the ICAPM: Evidence from the European stock market, *European Financial Management* 21, 1014-1052.

Maio, P., Santa-Clara, P., 2012. Multifactor models and their consistency with the ICAPM. *Journal of Financial Economics* 106, 586-613.

Merton, R.C, 1973. An intertemporal capital asset pricing model. *Econometrica* 41, 867-887.

Mossin, J., 1966, Equilibrium in a Capital Asset Market, *Econometrica* 34, 768–783.

Muir, T., and Moreira, A., 2017, Volatility Managed Portfolios, *Journal of Finance* 72, 1540-6261.

Nagel, S., Singleton, K., 2011, Estimation and Evaluation of Conditional Asset Pricing Models, *Journal of Finance* 66, 873-909.

Novy-Marx, R., 2015, Fundamentally, momentum is fundamental momentum, Working Paper University of Rochester.

Petkova, R., 2006. Do the Fama-French factors proxy for innovations in predictive variables? *Journal of Finance* 61, 581-612.

Roll, R., 1977. A critique of the asset pricing theory's tests Part I: On past and potential testability of the theory. *Journal of Financial Economics* 4, 129-176.

Rossi, A., Timmermann, A., 2015, Modeling Covariance Risk in Mertons ICAPM, *Review of Financial Studies* 28, 1428-1461.

Scruggs, J. T., 1998, Resolving the puzzling intertemporal relation between the market risk premium and conditional market variance: A two-factor approach, *Journal of Finance* 53, 575603.

Sharpe, W.F., 1964, Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk, *Journal of Finance* 19, 425-442.

Vasicek, O.A., 1973. A note on using cross-sectional information in Bayesian estimation of security betas. *Journal of Finance* 28, 1233-1239.

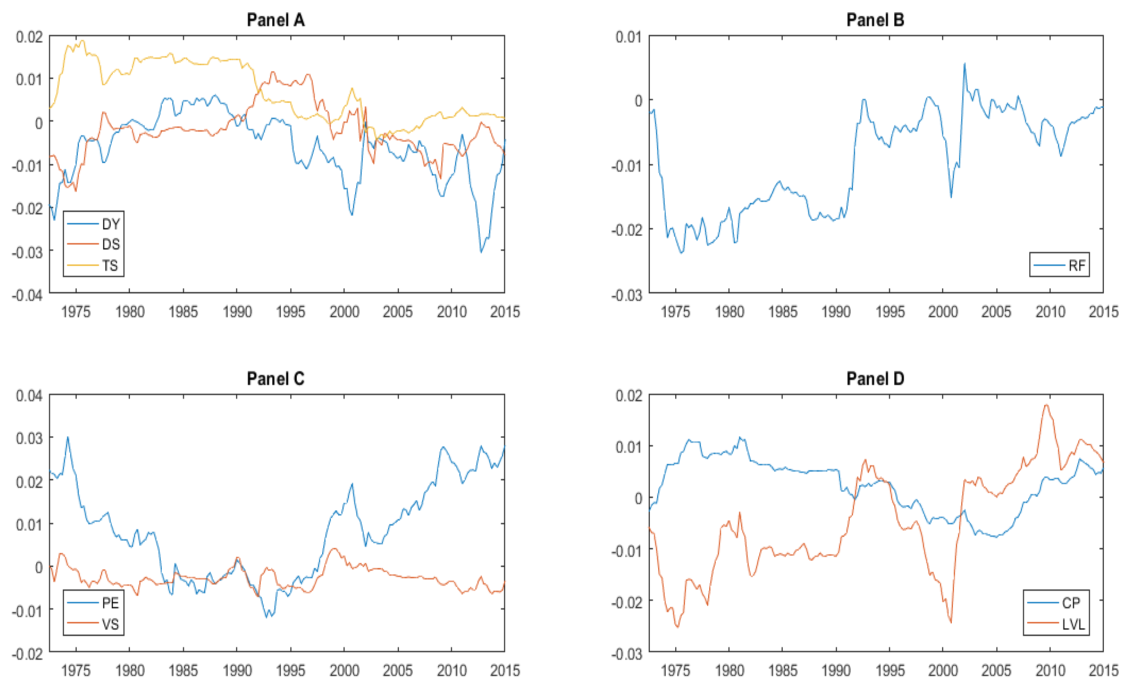

Figure 1: Rolling coefficients of consumption growth on lagged state variables
This figure presents rolling ten-year coefficient estimates from a regression of quarterly consumption growth on the state variables (lagged by one quarter) of a particular model, denoted b^1_{t,z_k} in the paper. We consider four models with multiple state variables: $\boldsymbol{z}_t = [DY, DS, TS]$, $\boldsymbol{z}_t = [DY, DS, RF]$, $\boldsymbol{z}_t = [TS, PE, VS]$, and $\boldsymbol{z}_t = [CP, LVL]$. For brevity, we present only the coefficient estimates for DY, DS, and TS from Model 1, RF from Model 2, PE and VS from Model 3, and CP and LVL from Model 4.

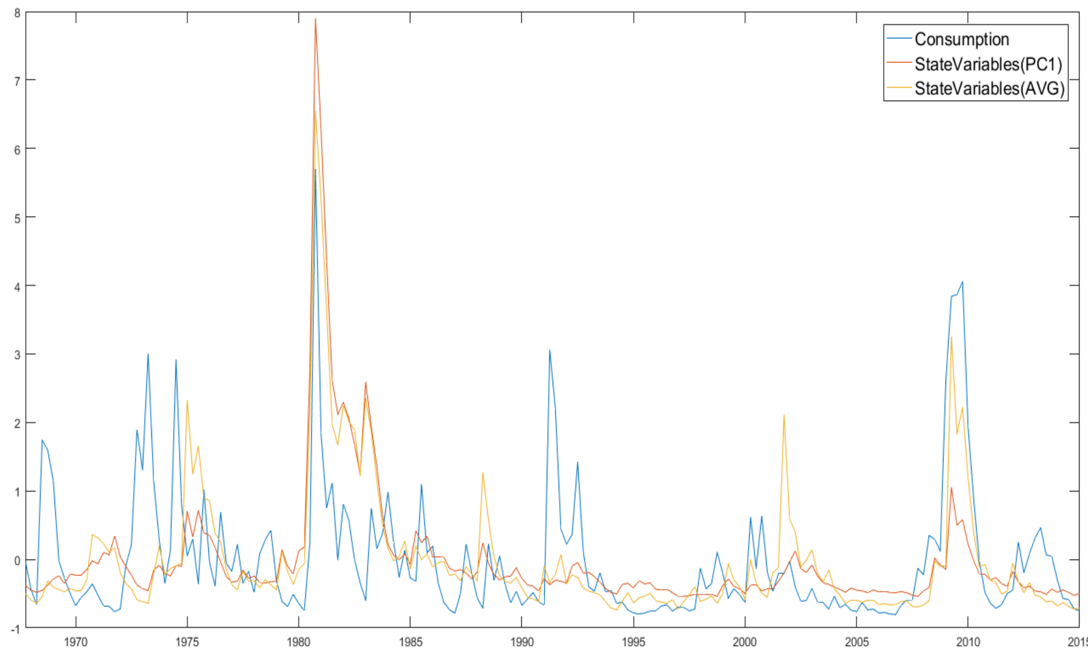

Figure 2: Conditional variance of the state variables and consumption growth
This figure plots the variance of consumption growth conditional on the state variables (Consumption, denoted $\sigma_{t,1}^{2,***}$ in Section 5) versus the conditional variance of the eight state variables (σ_{t,z_k}^2). We aggregate the σ_{t,z_k}^2 by taking an average across state variables (StateVariables(AVG)) and by estimating the first principal component (StateVariables(PC1)). All variance series are standardized to have mean zero and variance equal to one to facilitate interpretation.

Table I: Unconditional consumption growth predictability
This table presents an unconditional regression of future log consumption growth (cumulated over horizons of one and four quarters ($H = 1, 4$ in Panel A and B, respectively) on the state variables, $\boldsymbol{z}_t$, from four different models: $c_{t+1:t+K} = b_0 + \boldsymbol{b}'\boldsymbol{z}_t + e^u_{t+1:t+H}$. Model 1 sets $\boldsymbol{z}_t = [DY, DS, TS]$, Model 2 sets $\boldsymbol{z}_t = [DY, DS, RF]$, Model 3 sets $\boldsymbol{z}_t = [TS, PE, VS]$, and Model 4 sets $\boldsymbol{z}_t = [CP, LVL]$. t-statistics are presented underneath each coefficient estimate in brackets and calculated using Newey-West standard errors with H lags. The sample period runs from the second quarter of 1962 to the fourth quarter of 2014.

Model 1: DY, DS, TS		Model 2: DY, DS, RF		Model 3: TS, PE, VS		Model 4: CP, LVL	
Panel A: Quarterly growth ($H = 1$)							
DY	0.41	DY	0.27	TS	0.05	CP	0.27
	(2.41)		(1.40)		(0.26)		(1.35)
DS	-0.65	DS	-0.56	PE	0.11	LVL	0.13
	(-3.45)		(-2.99)		(0.70)		(0.70)
TS	0.25	RF	0.03	VS	-0.18		
	(1.46)		(0.16)		(-1.26)		
R^2	0.08	R^2	0.07	R^2	-0.01	R^2	0.02
Panel B: Annual growth ($H = 4$)							
DY	0.48	DY	0.41	TS	0.05	CP	0.29
	(2.28)		(1.81)		(0.29)		(2.05)
DS	-0.52	DS	-0.42	PE	-0.05	LVL	0.17
	(-2.43)		(-1.93)		(-0.24)		(0.81)
TS	0.27	RF	-0.08	VS	-0.04		
	(1.66)		(-0.33)		(-0.24)		
R^2	0.10	R^2	0.07	R^2	-0.01	R^2	0.05

Table II: Conditional consumption growth predictability

This table presents the results from a two-stage test for time-varying consumption growth predictability. In the first stage, we regress consumption growth on the lagged state variables over a backward-looking rolling window of 10 years (see Eq. (6) in the paper). In the second stage, we use the estimated rolling coefficients and the state variables observed at time t, i.e., $\widehat{a_{t,z}^{H}} + \widehat{\boldsymbol{b}_{t,z}^{H\prime}}\boldsymbol{z}_t$, to predict consumption growth from month $t+1$ to $t+H$ (see Eq. (5) in the paper). This setup ensures that we use no forward-looking information when we predict consumption growth in the second stage. We perform this test for each of the four models with $\boldsymbol{z}_t = [DY, DS, TS]$, $\boldsymbol{z}_t = [DY, DS, RF]$, $\boldsymbol{z}_t = [TS, PE, VS]$, and finally, $\boldsymbol{z}_t = [CP, LVL]$. The bootstrapped t-statistics use standard errors that are derived from 500 block-bootstrapped coefficient estimates. We report three measures of fit: the simple R^2, an R^2 comparison with the unconditional model of Table I, which is calculated as: $1 - var(e^c_{t+1:t+H})/var(e^u_{t+1:t+H})$, and, finally, an R^2 comparison with a conditional model that leaves out the state variables, such that consumption growth is only conditioned on its historical mean. The sample in the second stage regression runs from the second quarter of 1967 to the fourth quarter of 2014.

	Model 1	Model 2	Model 3	Model 4
Panel A: Quarterly growth ($H = 1$)				
d_0^1	0.01	0.01	0.01	0.01
	[1.80]	[2.28]	[1.57]	[1.32]
d_1^1	0.65	0.60	0.74	0.65
	[4.59]	[4.12]	[5.52]	[3.58]
R^2	0.29	0.26	0.29	0.16
R^2 Cond. vs Uncond.	0.25	0.23	0.28	0.10
R^2 Cond. vs Historical mean	0.26	0.22	0.26	0.12
Panel B: Annual growth ($H = 4$)				
d_0^4	0.01	0.01	0.01	0.01
	[2.24]	[2.77]	[2.44]	[1.90]
d_1^4	0.46	0.44	0.47	0.45
	[3.14]	[3.12]	[3.01]	[2.31]
R^2	0.21	0.23	0.17	0.11
R^2 Cond. vs Uncond.	0.15	0.21	0.14	-0.05
R^2 Cond. vs Historical mean	0.18	0.20	0.14	0.07

Table III: State variable risk premiums

This table presents the eight state variable risk premiums that will be the test assets for our conditional asset pricing tests. We estimate the state variable risk premiums with Fama and MacBeth (1973) cross-sectional regressions of quarterly individual stock returns on lagged exposures to the innovations in the state variables of a particular model (see Section 2.1). Panel A presents unconditional performance statistics for the risk premiums. Panel B regresses the risk premiums on a constant and a dummy variable indicating whether the lagged predictive coefficient for consumption growth, b^4_{t,z_k}, is above the sample median for that state variable: $\lambda_{t+1:t+4,z_k} = h_0 + h_1 I_{b^4_{t,z_k}>\text{med}} + e_{t+1:t+4,z_k}$. We report the coefficient estimates with corresponding t-statistics (based on Newey-West standard errors with 4 lags) underneath each estimate in parentheses. The sample period runs from the second quarter of 1967 to the last quarter of 2014.

	DY	DS	TS	RF	PE	VS	CP	LVL
Panel A: Unconditional state variable risk premiums								
Avg. Ret.	0.52	-6.22	5.67	-3.56	1.31	-5.08	3.74	1.91
	(0.26)	(-2.71)	(3.32)	(-1.84)	(0.68)	(-2.63)	(2.10)	(0.81)
Sharpe	0.04	-0.39	0.48	-0.27	0.10	-0.38	0.30	0.12
Panel B: Time-variation in risk premiums								
h_0	-2.56	-9.08	2.48	-0.89	-8.02	-6.24	0.43	0.48
	(-0.99)	(-2.84)	(2.05)	(-0.40)	(-2.43)	(-1.71)	(0.34)	(0.31)
h_1	8.65	6.73	6.42	3.46	2.61	6.39	2.86	1.57
	(2.52)	(2.00)	(1.72)	(0.96)	(0.75)	(1.72)	(0.80)	(0.43)

Table IV: Time-varying state variable risk premiums

This table presents the results from pooled predictive regressions of risk premiums on the lagged relation between the state variables and consumption growth: $\lambda_{t+1:t+H,z_k} = g_0 + g_1 b^H_{t,z_k} + e_{t+1:t+H,z_k}$ (see Section 4). Specifications 1 to 4 run the regressions for all eight state variables. To take care of the noise in the estimated b^H_{t,z_k}'s, we start by constructing dummy variables either indicating whether b^H_{t,z_k} is positive ($I_{b^H_{t,z_k}>0}$) in specification 1 and whether b^H_{t,z_k} is above its sample median in specification 2 ($I_{b^H_{t,z_k}>\text{med}}$). Specifications 3 and 4 use the raw measure, b^H_{t,z_k}, as independent variable, but specification 4 also includes state variable fixed effects. In specifications 5.1 to 5.4 we run the regression for the dummy specification with $I_{b^H_{t,z_k}>\text{med}}$ only for those state variables in a particular model: $\boldsymbol{z}_t = [DY, DS, TS]$ in 5.1, $\boldsymbol{z}_t = [DY, DS, RF]$ in Model 2, $\boldsymbol{z}_t = [TS, PE, VS]$ in Model 3, and $\boldsymbol{z}_t = [CP, LVL]$ in Model 4. The t-statistic in parentheses underneath each estimate is based on asymptotic Driscoll and Kraay (1998) standard errors; the t-statistic in brackets is based on standard errors derived from 500 block-bootstrapped coefficient estimates. The sample period runs from the second quarter of 1967 to the last quarter of 2014.

Specification	1	2	3	4	5.1	5.2	5.3	5.4
Quarterly returns ($H=1$)								
g_0	-2.18	-2.67	0.04		-3.09	-6.19	-2.09	1.26
	(-2.16)	(-2.22)	(0.08)		(-1.59)	(-2.62)	(-1.77)	(0.59)
	[-1.68]	[-2.60]	[0.07]		[-1.59]	[-3.19]	[-1.48]	[0.64]
g_1: $I_{b^1_{t,z_k}>0}$	4.90							
	(2.28)							
	[2.18]							
g_1: $I_{b^1_{t,z_k}>\text{med}}$		4.96			6.51	5.44	3.05	3.19
		(2.20)			(2.39)	(1.94)	(1.37)	(0.94)
		[2.88]			[2.72]	[2.14]	[1.41]	[0.98]
g_1: b^1_{t,z_k}			247.47	216.24				
			(2.06)	(1.87)				
			[2.05]	[1.97]				
$R^2(\times 100)$	0.74	0.79	0.83	2.25	1.33	0.89	0.35	0.30
Annual returns ($H=4$)								
g_0	-2.87	-2.50	0.31		-3.06	-5.05	-1.86	0.71
	(-2.71)	(-2.44)	(0.59)		(-1.92)	(-2.82)	(-1.96)	(0.43)
	[-2.19]	[-2.32]	[0.52]		[-1.58]	[-2.81]	[-1.16]	[0.35]
g_1: $I_{b^4_{t,z_k}>0}$	6.39							
	(2.90)							
	[2.83]							
g_1: $I_{b^4_{t,z_k}>\text{med}}$		5.17			7.27	4.15	2.38	4.42
		(2.54)			(2.76)	(1.99)	(1.10)	(1.43)
		[2.82]			[2.88]	[1.68]	[0.95]	[1.32]
g_1: b^4_{t,z_k}			269.77	238.52				
			(2.18)	(3.52)				
			[1.80]	[1.72]				
$R^2(\times 100)$	4.13	2.71	2.70	7.61	5.42	1.94	0.60	2.09
Fixed effects?	No	No	No	Yes	No	No	No	No

Table V: Simulation

We simulate state variables, consumption growth, and state variable risk premiums in an unconditional and a conditional ICAPM world. A detailed explanation of the simulations is in Appendix B. In the unconditional ICAPM, the relation between consumption growth and the lagged state variables is constant over time and so are the state variable risk premiums. In the conditional ICAPM, the relation between consumption growth and the state variables follows an auto-regressive processes, and state variable risk premiums are generated using this time-varying relation. All parameterizations are consistent with Model 1 with three state variables: $\boldsymbol{z}_t = [DY, DS, TS]$. Using these two simulated data sets, we compute the two-stage test for time-varying consumption growth predictability and report the R^2 comparison of the conditional and unconditional model (see Table II). In addition, we compute the unconditional average risk premium in the simulations in analogy to Table III. Finally, we conduct pooled predictive regressions as in Table IV by combining the state variable risk premiums from three separate simulations (using each simulation only once), so that we have a total of nine risk premiums in the pool (which is relative to eight in the data). Here, the table reports the slope coefficient for the specification 2 in which risk premiums are regressed on $I_{b^H_{t,z_k}>\text{med}}$, and the slope coefficient in specification 4 which uses the raw measure b^H_{t,z_k} as an independent variable and state variable fixed effects. The table reports the 5, 50 and 95 percentiles across simulations, and the fraction of times the coefficients was positive (> 0) and significantly positive at the 5% significance level ($>^* 0$).

	Percentiles			Percent of replications	
	5	50	95	> 0	$>^* 0$
Panel A: Unconditional ICAPM world					
R^2 Cond. vs Uncond.	-0.16	-0.10	-0.05	0.00	
Avg. Ret.	1.56	4.88	8.16	0.99	0.78
g_1: $I_{b^4_{t,z_k}>\text{med}}$ (Spec. 2)	-2.23	0.05	2.23	0.51	0.05
g_1: b^4_{t,z_k} (Spec. 4)	-230.66	4.66	230.83	0.51	0.05
Panel B: Conditional ICAPM world					
R^2 Cond. vs Uncond.	-0.10	0.02	0.20	0.60	
Avg. Ret.	-3.99	4.89	13.72	0.82	0.61
g_1: $I_{b^4_{t,z_k}>\text{med}}$ (Spec. 2)	0.58	3.91	7.47	0.97	0.78
g_1: b^4_{t,z_k} (Spec. 4)	92.67	387.71	694.87	0.98	0.86

Table VI: Robustness checks (I)

This table presents robustness checks focusing on the pooled regression of the eight state variable risk premiums on the dummy variable indicating whether the relation with consumption growth is above or below median for $H = 1, 4$: $\lambda_{t+1:t+H,z_k} = g_0 + g_1 I_{b^H_{t,z_k}>\text{med}} + e_{t+1:t+H,z_k}$. In Panel A we substitute measures of macroeconomic activity for consumption, such that b^H_{t,z_k} is estimated as the conditional predictive coefficient in a regression of real gross domestic product (GDP), industrial production growth (IPG), and the first principal component of the three macro growth series (PCRA) on the lagged state variables. Panel B replaces the Fama and MacBeth (1973) risk premium estimates at the stock-level with equal- or value-weighted high-minus-low quintile risk-sorted portfolio (HLEW and HLMV). The risk-sorted portfolios are sorted on exposures ($\boldsymbol{\delta}_{t,i,z}$'s) to DY, DS, and TS (from Model 1), RF (from Model 2), PE and VS (from Model 3), and CP and LVL (from Model 4). Next, we estimate the state variable risk premiums at the portfolio-level. We use a total of $8 \times 5 = 40$ value-weighted risk-sorted quintile portfolios ($40P$) to run the cross-sectional regression at each point in time to estimate conditional state variable risk premiums. In Panel C, we split the sample in two halves around the fourth quarter of 1990. For this exercise, we define the dummy indicator using the median over that particular sample half. In Panel D, we change the length of the rolling window used to estimate the relation between the state variable and future consumption growth to five and fifteen years, respectively. The t-statistics underneath each estimate are based on asymptotic standard errors calculated following Driscoll and Kraay (1998).

	Panel A: Macro-measures			Panel B: Portfolio returns			Panel C: Subsamples		Panel D: Rolling window	
	GDP	IPG	PCRA	HLEW	HLMV	$40P$	Pre-1990	Post-1990	5 year	15 year
				Quarterly returns ($H = 1$)						
g_0	-1.82	-1.86	-2.42	-2.55	-1.79	-2.13	-2.40	-1.51	-1.53	-1.41
	(-1.85)	(-1.64)	(-2.10)	(-2.83)	(-1.52)	(-2.17)	(-1.74)	(-1.15)	(-1.54)	(-1.23)
g_1: $I_{b^1_{t,z_k}>\text{med}}$	3.27	3.35	4.46	4.17	4.21	3.44	3.94	3.03	2.69	2.44
	(1.83)	(1.64)	(2.12)	(2.64)	(2.46)	(2.06)	(1.66)	(1.18)	(1.59)	(1.18)
				Annual returns ($H = 4$)						
g_0	-2.18	-1.57	-1.91	-1.88	-0.92	-1.56	-3.21	-1.52	-2.13	-2.10
	(-2.12)	(-1.49)	(-1.92)	(-2.33)	(-0.78)	(-1.56)	(-2.52)	(-1.49)	(-2.74)	(-1.78)
g_1: $I_{b^4_{t,z_k}>\text{med}}$	4.53	3.30	3.98	3.09	2.71	2.66	6.36	3.69	4.43	4.37
	(2.17)	(1.59)	(2.04)	(2.28)	(1.83)	(1.49)	(2.32)	(2.00)	(2.95)	(1.97)

Table VII: Time-varying risk premiums controlling for benchmark factors
This table analyzes whether our conslusions are robust to controlling for contemporaneous exposure of state variable risk premiums to the benchmark traded factors of the CAPM as well as the Fama-French three- and five-factor models. To this end, we run a time series regression of the state variable risk premiums on the contemporaneous factor returns and perform our analyses on the estimated intercept plus residual. In Panel A we run the predictive regression of orthogonalized risk premiums on the lagged relation between the state variable and consumption growth (see Table IV). In Panel B we regress orthogonalized risk premiums on the conditional relation between each state variable and future consumption growth controlling for the GARCH(1,1) conditional variance of the state variables (see Table VIII). In Panel C we regress orthogonalized risk premiums on the conditional relation between each state variable and future consumption growth controlling for macroeconomic uncertainty, as measured by the conditional variance of consumption growth (see Table X). In all cases, we present results for the annual frequency and for the dummy specifications. In Panels B and C, we focus on the predicted risk premiums in quarters with either positive or negative b^H_{t,z_k} and conditional variance below or above median. The t-statistics in parenthesis are based on asymptotic standard errors calculated following Driscoll and Kraay (1998).

Panel A: Conditional relation with consumption ($I_{b^4_{t,z_k}>\text{med}}$)			
λ_{t+1,z_k}'s orthogonalized from	CAPM	FF3M	FF5M
g_0	-2.89	-2.54	-2.19
	(-2.79)	(-2.61)	(-2.30)
g_1	5.07	4.87	4.71
	(2.45)	(2.43)	(2.34)
$R^2 \times 100$	2.71	2.89	2.75

Panel B: Conditional variance ($I_{b^4_{t,z_k}>0}, I_{\sigma^2_{t,z_k}>\text{med}}$)				Panel C: Macro-uncertainty ($I_{b^4_{t,z_k}>0}, I_{\sigma^{2,*}_{t,z,4}>\text{med}}$)			
λ_{t+1,z_k}'s orthogonalized from	CAPM	FF3M	FF5M		CAPM	FF3M	FF5M
Low b^4_{t,z_k}, High σ^2_{t,z_k}	-6.37	-4.63	-4.46	Low b^4_{t,z_k}, High $\sigma^{2,*}_{t,z,4}$	-4.59	-3.68	-3.36
	(-3.94)	(-3.14)	(-2.92)		(-3.37)	(-2.92)	(-2.58)
Low b^4_{t,z_k}, Avg. σ^2_{t,z_k}	-0.98	-0.50	0.12	Low b^4_{t,z_k}, Avg. $\sigma^{2,*}_{t,z,4}$	-2.41	-1.19	-0.69
	(-1.28)	(-0.66)	(0.18)		(-2.31)	(-1.20)	(-0.71)
High b^4_{t,z_k}, Avg. σ^2_{t,z_k}	-0.03	-0.70	-0.75	High b^4_{t,z_k}, Avg. $\sigma^{2,*}_{t,z,4}$	1.68	0.76	0.75
	(-0.03)	(-0.84)	(-0.88)		(1.43)	(0.67)	(0.64)
High b^4_{t,z_k}, High σ^2_{t,z_k}	6.12	5.41	5.66	High b^4_{t,z_k}, High $\sigma^{2,*}_{t,z,4}$	4.86	4.40	4.63
	(3.18)	(2.90)	(2.94)		(2.55)	(2.36)	(2.42)

Table VIII: Conditional variance of state variables and risk premiums

This table presents pooled predictive regressions of state variable risk premiums on the conditional relation between each state variable and future consumption growth controlling for the GARCH(1,1) conditional variance of the state variables (see Section 3.1): $\lambda_{t+1:t+H,z_k} = g_0 + g_1 b^H_{t,z_k} + g_2 \sigma^2_{t,z_k} + g_3 \sigma^2_{t,z_k} \times b^H_{t,z_k} + e_{t+1:t+H,z_k}$. The dependent variables are the risk premiums for the eight state variables of interest (i.e., for DY, DS, and TS in Model 1, RF in Model 2, PE and VS in Model 3, and CP and LVL in Model 4). In the right block of results, we also estimate a dummy-specification in which we replace b^H_{t,z_k} with the indicator $I_{b^H_{t,z_k}>0}$ and σ_{t,z_k} with the indicator $I_{\sigma^2_{t,z_k}>\text{med}}$. We present the coefficient estimates for the annual horizon $H = 4$ in Panel A. To grasp the economic significance of the estimates, Panel B presents the predicted risk premium in quarters with b^H_{t,z_k} one standard below or above the mean versus σ^2_{t,z_k} at the mean or one standard deviation above the mean. For the dummy specification, we present the implied risk premium in quarters with positive or negative b^H_{t,z_k} and conditional variance below or above median. The t-statistic in parentheses underneath each estimate is based on asymptotic Driscoll and Kraay (1998) standard errors; the t-statistic in brackets is based on standard errors derived from 500 block-bootstrapped coefficient estimates.

Continuous: b^4_{t,z_k} and σ^2_{t,z_k}		Dummy: $I_{b^4_{t,z_k}>0}$ and $I_{\sigma^2_{t,z_k}>\text{med}}$	
Panel A: Coefficient estimates			
g_0	0.22	g_0	-0.27
	(0.44)		(-0.32)
	[0.37]		[-0.20]
g_1: b^4_{t,z_k}	214.01	g_1: $I_{b^4_{t,z_k}>0}$	0.39
	(2.02)		(0.30)
	[1.41]		[0.19]
g_2: σ^2_{t,z_k}	-0.08	g_2: $I_{\sigma^2_{t,z_k}>\text{med}}$	-5.58
	(-0.10)		(-3.28)
	[-0.09]		[-2.89]
g_3: $b^4_{t,z_k} \times \sigma^2_{t,z_k}$	225.94	g_3: $I_{b^4_{t,z_k}>0} \times I_{\sigma^2_{t,z_k}>\text{med}}$	11.88
	(2.63)		(3.59)
	[1.39]		[4.04]
$R^2(\times 100)$	6.29		7.68
Panel B: Predicted state variable risk premiums			
Low b^4_{t,z_k}, High σ^2_{t,z_k}	-4.69		-5.85
	(-2.62)		(-3.54)
	[-2.11]		[-3.17]
Low b^4_{t,z_k}, Avg. σ^2_{t,z_k}	-2.13		-0.27
	(-1.92)		(-0.32)
	[-1.68]		[-0.20]
High b^4_{t,z_k}, Avg. σ^2_{t,z_k}	2.18		0.12
	(1.89)		(0.12)
	[1.63]		[0.09]
High b^4_{t,z_k}, High σ^2_{t,z_k}	4.17		6.42
	(2.29)		(3.37)
	[1.85]		[3.48]

Table IX: Robustness checks (II)

This table presents robustness checks for Tables VIII (Panel A), X (Panel B), and XI (Panel C), focusing on the predicted risk premiums. These premiums follow from pooled regressions (at the annual horizon $H = 4$) for the eight state variable risk premiums on two dummy variables (and their interaction) indicating whether (i) the conditional relation between a state variable and consumption-investment opportunities is high or low and (ii) the conditional variance of state variables (Panels A and C) or consumption growth (Panel B) is high or low. The left three columns substitute measures of macroeconomic activity for consumption growth, such that we estimate b^H_{t,z_k} as the conditional predictive coefficient in a regression of real gross domestic product (GDP), industrial production growth (IPG), and the first principal component of the three macro growth series (PCRA) on the lagged state variables. The right three columns replace the Fama and MacBeth (1973) risk premium estimates at the stock-level with equal- or value-weighted high-minus-low quintile risk-sorted portfolio (HLEW and HLMV). The risk-sorted portfolios are sorted on exposures ($\boldsymbol{\delta}_{t,i,z}$'s) to DY, DS, and TS (from Model 1), RF (from Model 2), PE and VS (from Model 3) and CP and LVL (from Model 4). Finally, we estimate the state variable risk premiums at the portfolio-level. To do so, we use a total of $8 \times 5 = 40$ value-weighted risk-sorted quintile portfolios ($40P$) instead of individual stocks to run the cross-sectional regression at each point in time to estimate conditional state variable risk premiums. The t-statistics in parentheses are based on asymptotic standard errors calculated following Driscoll and Kraay (1998).

	GDP	IPG	PCRA	HLEW	HLMV	$40P$
Panel A: Conditional variance of the state variables						
Low b^H_{t,z_k}, High σ^2_{t,z_k}	-4.02	-3.04	-4.27	-5.00	-2.34	-1.75
	(-2.56)	(-1.67)	(-2.61)	(-4.00)	(-1.48)	(-1.67)
Low b^H_{t,z_k}, Avg. σ^2_{t,z_k}	-0.48	-0.48	-0.32	-1.24	-0.07	-0.44
	(-0.61)	(-0.52)	(-0.37)	(-1.97)	(-0.06)	(-0.50)
High b^H_{t,z_k}, Avg. σ^2_{t,z_k}	0.39	0.29	0.14	0.71	0.96	0.21
	(0.37)	(0.30)	(0.14)	(0.67)	(0.55)	(0.20)
High b^H_{t,z_k}, High σ^2_{t,z_k}	5.72	4.20	5.76	4.51	3.37	1.14
	(2.55)	(1.97)	(2.63)	(4.21)	(2.29)	(0.70)
Panel B: Macroeconomic uncertainty						
Low b^4_{t,z_k}, High $\sigma^{2,*}_{t,z,4}$	-3.66	-2.91	-3.69	-4.53	-2.41	-1.38
	(-2.62)	(-2.55)	(-2.72)	(-4.78)	(-1.58)	(-1.35)
Low b^4_{t,z_k}, Avg. $\sigma^{2,*}_{t,z,4}$	-0.88	-0.64	-0.88	-1.48	0.14	-0.73
	(-0.97)	(-0.36)	(-0.77)	(-2.10)	(0.12)	(-0.68)
High b^4_{t,z_k}, Avg. $\sigma^{2,*}_{t,z,4}$	2.48	1.55	2.48	1.98	1.36	-0.52
	(2.13)	(1.17)	(1.82)	(2.03)	(0.94)	(-0.43)
High b^4_{t,z_k}, High $\sigma^{2,*}_{t,z,4}$	3.64	2.85	3.34	3.52	3.15	1.93
	(1.71)	(1.53)	(1.55)	(2.88)	(1.90)	(1.26)
Panel C: Conditional variance of the state variables net of macro-uncertainty						
Low b^H_{t,z_k}, High σ^2_{t,z_k}	-0.07	-1.28	-1.00	-3.33	-1.32	-1.68
	(-0.05)	(-0.45)	(-0.63)	(-2.48)	(-0.79)	(-0.92)
Low b^H_{t,z_k}, Avg. σ^2_{t,z_k}	-2.89	-1.95	-2.71	-2.91	-1.08	-0.89
	(-2.79)	(-2.08)	(-2.51)	(-3.79)	(-0.88)	(-1.09)
High b^H_{t,z_k}, Avg. σ^2_{t,z_k}	2.71	2.18	2.46	2.33	1.81	1.14
	(1.74)	(1.61)	(1.63)	(2.33)	(1.33)	(0.88)
High b^H_{t,z_k}, High σ^2_{t,z_k}	4.26	2.23	4.33	4.23	3.84	-0.79
	(2.06)	(0.92)	(2.02)	(2.81)	(1.87)	(-0.40)

Table X: Macroeconomic uncertainty and state variable risk premiums
This table presents pooled predictive regressions of state variable risk premiums on the conditional relation between each state variable and future consumption growth controlling for macroeconomic uncertainty, as measured by the conditional variance of consumption growth: $\lambda_{t+1:t+H,z_k} = g_0 + g_1 b^H_{t,z_k} + g_2\sigma^2_{t,H} + g_3\sigma^2_{t,H} \times b^H_{t,z_k} + e_{t+1:t+H,z_k}$. $\sigma^2_{t,H}$ is one of the three uncertainty measures described in Section 5: $\sigma^{2,*}_{t,z,H}$, $\sigma^{2,**}_{t,z,H}$, and $\sigma^{2,***}_{t,H}$. The dependent variable returns are the risk premiums for eight state variables, as before. In the right block of results, we also estimate a dummy-specification where we replace b^H_{t,z_k} with the indicator $I_{b^H_{t,z_k}>0}$ and $\sigma_{t,H}$ with the indicator $I_{\sigma^2_{t,H}>\text{med}}$. We present the coefficient estimates for the annual horizon $H = 4$ in Panel A. To grasp the economic significance of the estimates, we present in Panel B the predicted risk premium in quarters with b^H_{t,z_k} one standard below or above the mean versus $\sigma^2_{t,H}$ at the mean or one standard deviation above the mean. For the dummy specification, we present the predicted risk premium in quarters with positive or negative b^H_{t,z_k} and conditional variance below or above median. The t-statistic in parentheses underneath each estimate is based on asymptotic Driscoll and Kraay (1998) standard errors; the t-statistic in brackets is based on standard errors derived from 500 block-bootstrapped coefficient estimates.

	Continuous: b^4_{t,z_k} and $\sigma^2_{t,4}$				Dummy: $I_{b^4_{t,z_k}>0}$ and $I_{\sigma^2_{t,4}>\text{med}}$		
	$\sigma^{2,*}_{t,z,4}$	$\sigma^{2,**}_{t,z,4}$	$\sigma^{2,***}_{t,4}$		$\sigma^{2,*}_{t,z,4}$	$\sigma^{2,**}_{t,z,4}$	$\sigma^{2,***}_{t,4}$
Panel A: Coefficient estimates							
g_0	0.24	0.41	0.11	g_0	-1.56	-1.65	-1.81
	(0.41)	(0.58)	(0.21)		(-1.46)	(-1.53)	(-1.62)
	[0.35]	[0.61]	[0.16]		[-1.25]	[-1.21]	[-1.44]
g_1: b^4_{t,z_k}	64.64	170.08	60.04	g_1: $I_{b^4_{t,z_k}>0}$	3.57	3.47	3.64
	(0.48)	(1.40)	(0.46)		(1.78)	(1.75)	(1.79)
	[0.36]	[1.09]	[0.34]		[1.66]	[1.60]	[1.68]
g_2: $\sigma^2_{t,4}$	0.01	-0.02	0.09	g_2: $I_{\sigma^2_{t,4}>\text{med}}$	-2.64	-2.51	-2.18
	(0.06)	(-0.29)	(0.73)		(-1.77)	(-1.82)	(-1.36)
	[0.05]	[-0.17]	[0.41]		[-1.68]	[-1.47]	[-1.37]
g_3: $b^4_{t,z_k} \times \sigma^2_{t,4}$	121.23	11.48	92.51	g_3: $I_{b^4_{t,z_k}>0} \times I_{\sigma^2_{t,4}>\text{med}}$	5.64	5.81	5.47
	(2.35)	(1.43)	(2.39)		(1.83)	(2.24)	(1.60)
	[1.33]	[0.66]	[1.01]		[2.14]	[2.14]	[2.00]
$R^2(\times 100)$	5.01	3.36	5.36		4.93	4.98	4.90
Panel B: Predicted state variable risk premiums							
Low b^4_{t,z_k}, High $\sigma^2_{t,4}$	-4.81	-3.47	-4.66		-4.20	-4.16	-3.99
	(-3.44)	(-1.93)	(-3.06)		(-2.92)	(-3.00)	(-2.72)
	[-2.29]	[-2.54]	[-2.13]		[-2.39]	[-2.50]	[-2.25]
Low b^4_{t,z_k}, Avg. $\sigma^2_{t,4}$	-2.40	-2.26	-2.30		-1.56	-1.65	-1.81
	(-2.22)	(-1.99)	(-2.09)		(-1.46)	(-1.53)	(-1.62)
	[-1.96]	[-2.06]	[-1.88]		[-1.25]	[-1.21]	[-1.44]
High b^4_{t,z_k}, Avg. $\sigma^2_{t,4}$	2.49	2.42	2.49		2.01	1.82	1.83
	(2.00)	(1.86)	(1.98)		(1.72)	(1.43)	(1.64)
	[1.91]	[1.94]	[1.91]		[1.68]	[1.51]	[1.52]
High b^4_{t,z_k}, High $\sigma^2_{t,4}$	4.56	3.12	4.94		5.00	5.11	5.12
	(2.64)	(2.34)	(2.73)		(2.57)	(2.75)	(2.52)
	[2.03]	[2.11]	[2.15]		[3.01]	[3.14]	[3.04]

Table XI: Conditional variance of state variables net of macroeconomic uncertainty

This table is similar to Table VIII and presents predicted risk premiums estimated from a pooled regression of state variable risk premiums on the conditional relation between each state variable and future consumption growth (b^H_{t,z_k}) as well as the conditional variance of the state variables (σ^2_{t,z_k}). In this case, the conditional variance measures control for quarters where macroeconomic uncertainty ($\sigma^2_{t,z,H}$) is high. To do so, we set σ^2_{t,z_k} equal to its time-series median in quarters when macro uncertainty $\sigma^2_{t,z,H}$ is above median. In the dummy specification, we set $I_{\sigma^2_{t,z_k}>\text{med}}$ to zero in quarters when $\sigma^2_{t,z,H}$ is also above median. To proxy for macroeconomic uncertainty, we consider each of the three measures of the conditional variance of consumption growth ($\sigma^{2,*}_{t,z,H}, \sigma^{2,**}_{t,z,H}, \sigma^{2,***}_{t,z,H}$). The predicted risk premiums are for the four cases also analyzed in Table VIII (Low b^H_{t,z_k}, High σ^2_{t,z_k}; Low b^H_{t,z_k}, Low σ^2_{t,z_k}; High b^H_{t,z_k}, Low σ^2_{t,z_k}; High b^H_{t,z_k}, High σ^2_{t,z_k}) and at the annual horizon ($H = 4$). The t-statistic in parentheses underneath each estimate is based on asymptotic Driscoll and Kraay (1998) standard errors; the t-statistic in brackets is based on standard errors derived from 500 block-bootstrapped coefficient estimates.

	Continuous: b^4_{t,z_k} and σ^2_{t,z_k}			Dummy: $I_{b^4_{t,z_k}>0}$ and $I_{\sigma^2_{t,z_k}>\text{med}}$		
	$\sigma^{2,*}_{t,z,4}$	$\sigma^{2,**}_{t,z,4}$	$\sigma^{2,***}_{t,4}$	$\sigma^{2,*}_{t,z,4}$	$\sigma^{2,**}_{t,z,4}$	$\sigma^{2,***}_{t,4}$
	Predicted state variable risk premiums					
Low b^4_{t,z_k}, High σ^2_{t,z_k}	-2.89	-3.43	-3.84	-3.63	-4.17	-4.73
	(-3.71)	(-2.06)	(-3.45)	(-2.09)	(-2.57)	(-2.41)
	[-1.80]	[-2.08]	[-2.39]	[-2.01]	[-2.19]	[-2.61]
Low b^4_{t,z_k}, Avg. σ^2_{t,z_k}	-2.35	-2.50	-2.49	-2.68	-2.52	-2.39
	(-2.01)	(-1.93)	(-1.92)	(-2.45)	(-2.42)	(-2.21)
	[-1.95]	[-2.28]	[-2.05]	[-1.93]	[-1.90]	[-1.72]
High b^4_{t,z_k}, Avg. σ^2_{t,z_k}	2.56	2.43	2.62	3.14	2.94	3.18
	(1.88)	(1.91)	(1.82)	(2.23)	(2.11)	(2.16)
	[1.99]	[1.95]	[2.02]	[2.43]	[2.28]	[2.46]
High b^4_{t,z_k}, High σ^2_{t,z_k}	3.30	4.07	2.14	4.84	5.44	4.67
	(3.12)	(2.11)	(3.15)	(2.27)	(2.49)	(2.24)
	[1.89]	[2.32]	[1.20]	[2.76]	[3.12]	[2.62]