Good and Bad Variance Premia and Expected Returns *

Mete Kilic
University of Southern California

Ivan Shaliastovich
University of Wisconsin-Madison

July 5, 2017

Abstract

We measure "good" and "bad" variance premia that capture risk compensations for the realized variation in positive and negative market returns, respectively. The two variance premium components jointly predict excess returns over the next 1 and 2 years with statistically significant positive (negative) coefficients on the good (bad) component. The R^2s reach about 10% for aggregate equity and portfolio returns, and 20% for corporate bond returns. To explain the new empirical evidence, we develop a model that highlights the differential impact of upside and downside risk on equity and variance risk premia.

*First draft: February 2015. Mete Kilic: mkilic@marshall.usc.edu, Department of Finance and Business Economics, Marshall School of Business, University of Southern California; Ivan Shaliastovich: ivan.shaliastovich@wisc.edu, Finance Department, Wisconsin School of Business, University of Wisconsin-Madison. We are grateful to Torben Andersen, Geert Bekaert, Bjorn Eraker, Nicola Fusari, Juan M. Londono, Jessica Wachter, Amir Yaron, Hao Zhou, two anonymous referees and Neng Wang (department editor), as well as seminar participants at Emory University, Temple University, University of Maryland, University of North Carolina at Chapel Hill, University of Wisconsin-Madison, the Wharton School, the 2015 Midwest Finance Association meetings, 2016 American Finance Association meetings, and the JHU Conference on The Role of Derivatives in Asset Pricing for helpful comments. We thank The Rodney L. White Center for Financial Research for financial support.

1 Introduction

An important challenge in finance is to link expected excess asset returns to measures of the financial market risk. Using statistical estimates of equity return variance to measure risk, the early asset pricing literature has generally produced inconclusive evidence for the risk and return relation.[1] A more recent strand of this literature relies on the information in option prices to gauge time-varying risk compensations in the data. In particular, the variance premium, defined as the difference between the physical and risk-neutral expectations of return variation, has been shown to be a robust predictor of asset returns at maturities of 3 to 6 months. Because of its significant predictive power for short-term asset returns, the variance premium is often viewed as a volatility-related measure of transient risk in financial markets. In this paper, we show that there are two distinct notions of the variance premia, which represent compensations for "good" and "bad" variation in market returns. In the data, good and bad variance premia have opposite implications for the level and variation in the risk compensation in financial markets. We provide an economic model to explain the new empirical evidence, and highlight the difference in investors' risk attitudes towards upside and downside uncertainty risks.

Our variance premium measures are straightforward to compute and have significant implications for our understanding of investors' attitudes to variance risks. First, the good variance premium is positive, on average, while the bad variance premium is negative most of the time. This implies that investors dislike bad variance risk, and like the states with high good market variance. Second, the good and bad variance premia jointly predict excess equity market returns with statistically significant coefficients at the 1- and 2-year horizons with R^2 values of 7% and 9%, respectively. In contrast, the predictive power of the total variance premium at horizons longer than 6 months is essentially zero. Notably, both the good and bad components of the variance premium need to be included to obtain high return predictability: the predictive power significantly drops in univariate regressions using the individual components of the variance premium. This suggests that both upward and downward risks play an important role in capturing time variation

[1] See French, Schwert, and Stambaugh (1987), Campbell and Hentschel (1992), Glosten, Jagannathan, and Runkle (1993), and Whitelaw (1994).

in asset risk premia.[2] Finally, the predictive coefficient on the good variance premium is positive, while it is negative for the bad variance premium. Because the good and bad variance premia drive the total variance premium in the same direction but predicted returns in the opposite direction, the variance premium decomposition helps uncover long-horizon return predictability, absent when the total variance premium itself is used as a predictor. This pattern holds across various markets, including aggregate equity, corporate bonds, and the cross-section of equity portfolios.

To provide an explanation for our empirical findings, we start with an illustrative no-arbitrage framework with exogenously specified equity price and state-price density processes. The equity price and state-price densities are driven by upward ("good") and downward ("bad") jumps. Using comparative statics, we show that the equity premium is increasing in both good and bad jump intensities. At the same time, the good variance premium is increasing in good jump intensity, while the bad variance premium is decreasing in bad jump intensity. This can explain a negative coefficient for the bad variance premium, and positive coefficient for the good variance premium in return predictability regressions. Using the total variance premium conflates the two opposite effects of good and bad jump intensities, and thus masks the relation between the variance and the risk premium. We further show that the risk and return relationship is ambiguous using the good and bad volatilities themselves, in place of the good and bad variance premia. The two variance premia are able to isolate the two components related to jump risks, and thus provide a cleaner identification of the risk and return tradeoff in the financial markets. We further consider a general equilibrium extension of our framework which builds on the intuition of the illustrative model. The model allows us to provide an economic foundation for the assumptions made in the illustrative model and disciplines the signs and magnitudes of model parameters for a quantitative evaluation of the economic channels.

Related literature. Our paper is related to the literature that links the variance premium to uncertainty about economic fundamentals. Bollerslev, Tauchen, and Zhou (2009) show the strong

[2] See Ang, Chen, and Xing (2006) and Lettau, Maggiori, and Weber (2014), who emphasize the importance of the downside market risk for the cross-section of asset returns.

short-horizon predictive power of the variance premium for aggregate equity returns.[3] In their model, the variance premium is related to the time-varying volatility of conditional consumption variance.[4] Drechsler and Yaron (2011) extend the long-run risk model of Bansal and Yaron (2004) by allowing for multiple volatility factors and incorporating jumps in the expected growth and variance of consumption. The variance premium in this model becomes a proxy for the variance factor with the lowest persistence. Eraker (2008) studies the implications of volatility jumps for the variance premium. The empirical evidence in our paper suggests that the variance premium is driven by at least two factors that are related to long-horizon expected returns. Our economic model accounts for this two-factor structure by linking the components of the variance premium to the good and bad jump intensities.

Our paper is also related to the recent literature that highlights the variation in positive and negative shocks to fundamentals. Our proxies for the good and bad variance premia correspond to the differences between the good and bad realized and implied variance of returns. Good (bad) implied variance measures the conditional risk-neutral expectation of the 1-month squared positive (negative) log equity return, computed from option prices as in Bakshi and Madan (2000) and Bakshi, Kapadia, and Madan (2003). The good and bad realized variance measures are computed using the high-frequency return data, following Barndorff-Nielsen, Kinnebrock, and Shephard (2010). Using a similar econometric approach to identify realized good and bad variations, Segal, Shaliastovich, and Yaron (2015) study the implications of the good and bad uncertainty, measured from macroeconomic data, for real growth and asset valuations. Bekaert and Engstrom (2015) consider a habit formation model with distinct variation in the volatilities of positive and negative Gamma shocks to fundamentals to address several stylized asset pricing puzzles, including the predictability of returns by the total variance premium. Tsai and Wachter (2015) construct a model with distinct time variation in the probabilities of rare booms and disasters to explain the value premium. Our model integrates the economic channels used in this literature in a

[3]For long-horizon predictors of equity and corporate bond returns, see Keim and Stambaugh (1986), Campbell and Shiller (1988), Fama and French (1989), Cochrane (2008), among others.

[4]Short-horizon stock return predictability results using the variance premium have been extended since then. See Zhou (2009), Han and Zhou (2012), Bollerslev, Marrone, Xu, and Zhou (2014), and Bali and Zhou (2014), among others.

parsimonious way, focusing on the novel empirical evidence for a link between the good and bad variance premia and expected returns.

Our paper also contributes to a voluminous literature that shows that asset returns are predictable by volatility and jump risk factors.[5] Kelly and Jiang (2014) construct a tail risk measure using the cross-section of equity returns, and show that it contains information about expected excess equity returns. Bollerslev, Todorov, and Xu (2015) statistically disentangle the diffusive and jump components of the variance premium and show that this decomposition leads to stronger return predictability than previously shown, driven by the variation in the left jump tail. Guo, Wang, and Zhou (2014) find that realized positive and negative jump volatilities jointly predict short-horizon excess stock returns and economic fundamentals, while total jump variation has no significant predictive power. Feunou, Jahan-Parvar, and Okou (2014) study the empirical implications of decomposing the VIX index into the components computed using call and put options and find that most of the variance premium is related to downside risk. In our paper, we consider an alternative parsimonious decomposition of the variance premium into the components associated with "good" and "bad" events and provide an economic model to explain our empirical findings.

Finally, our paper is related to the recent literature that highlights the importance of separating upside and downside market volatilities. Feunou, Jahan-Parvar, and Tédongap (2013), Patton and Sheppard (2013), and Bekaert, Engstrom, and Ermolov (2015) identify fluctuations in these volatility components using historical return data and consider their implications for the dynamics of equity returns. In our paper, we utilize option data and construct variance premium measures, which we show can help to better isolate separate variations in the negative and positive components of the return distribution, related to the expected excess returns.

The paper is organized as follows. Section 2 provides an illustrative model to elucidate the economic mechanisms behind the variance premia. Section 3 describes the benchmark empirical analysis, along with several robustness tests. Section 4 summarizes the economic model and its quantitative implications. Section 5 concludes.

[5]The literature that highlights jumps in prices includes Bakshi, Cao, and Chen (1997), Duffie, Pan, and Singleton (2000), Pan (2002), Andersen, Benzoni, and Lund (2002), Eraker, Johannes, and Polson (2003), Eraker (2004), Broadie, Chernov, and Johannes (2007), Santa-Clara and Yan (2010).

2 An Illustrative Model

In this section, we present a simple *iid* model to illustrate the economic mechanism behind the variance premium, its decomposition into good and bad components, and their relation to the equity premium. The qualitative implications of this model, derived under comparative statics, motivate our subsequent empirical analysis in Section 3. In Section 4 we provide a general equilibrium intuition for the reduced-form assumptions of the illustrative model.

2.1 Model Setup

In our illustrative example, the asset price follows the process:

$$\frac{dS_t}{S_{t^-}} = \mu_S \, dt + \sigma \, dB_t + \left(e^{Z_{g,t}} - 1\right) \, dN_{g,t} + \left(e^{-Z_{b,t}} - 1\right) \, dN_{b,t}, \tag{1}$$

where μ_S is the drift, B_t is a standard Brownian motion, and $N_{g,t}$ and $N_{b,t}$ are Poisson processes with jump intensities λ_g and λ_b. The jump sizes $Z_{g,t}$ and $Z_{b,t}$ are *iid* over time, and take only non-negative values. As a result, $Z_{g,t}$ represents a "good" jump while $Z_{b,t}$ is a "bad" asset-price jump.[6] For parsimony, expected returns, jump intensities, and diffusive volatility are assumed to be constant so that the model is *iid*.

The state-price density π_t is driven by the same shocks that move the asset price, and follows the process:

$$\frac{d\pi_t}{\pi_{t^-}} = \mu_\pi \, dt - \Lambda\sigma \, dB_t + \left(e^{-\Lambda Z_{g,t}} - 1\right) \, dN_{g,t} + \left(e^{\Lambda Z_{b,t}} - 1\right) \, dN_{b,t}, \tag{2}$$

where $\Lambda > 0$ denotes the market price of risk. For parsimony, we assume that the magnitudes of the market prices of all risks are the same, in absolute value. The state-price density is a proxy

[6] Formally, $Z_{g,t}$ and $Z_{b,t}$ are jumps in the log asset price. By Ito's lemma,

$$d \log S_t = \mu_{S,t}^{\log} \, dt + \sigma dB_t + \sum_{i=g,b} Z_{i,t} \, dN_{i,t}$$

where $\mu_{S,t}^{\log} = \mu_{S,t} - \frac{1}{2}\sigma^2$.

for marginal utility, and high marginal utility states correspond to "bad" times under concave utility functions. Good asset-price jumps capture good economic states with low marginal utility; hence, good jumps have a positive market price of risk. On the other hand, bad asset-price jumps increase marginal utility, and have a negative market price of risk.[7]

Our model belongs to the affine jump-diffusion class, and can be solved using standard machinery. In particular, the dynamics of the asset price under the risk-neutral measure takes the same form as under the physical measure given in (1), but the distribution of good and bad jumps and their intensities have to be adjusted to take into account the compensation for jump risks Specifically, because good jumps represent good economic states, they arrive less frequently under the risk-neutral measure, and their distribution is tilted towards smaller values. On the other hand, bad jumps represent adverse states for the investor, so the bad jump intensity is higher, and larger bad jumps are more likely under the risk-neutral measure.

2.2 Risk Compensations

The instantaneous risk premium on the asset is constant, and is given by

$$\begin{aligned} r_t^e - r_t^f &= \Lambda\sigma^2 + \mathbb{E}\left[\left(e^{-\Lambda Z_{g,t}} - 1\right)\left(1 - e^{Z_{g,t}}\right)\right]\lambda_g + \mathbb{E}\left[\left(e^{\Lambda Z_{b,t}} - 1\right)\left(1 - e^{-Z_{b,t}}\right)\right]\lambda_b \\ &\equiv \Lambda\sigma^2 + \beta_g^{ep}\,\lambda_g + \beta_b^{ep}\,\lambda_b, \end{aligned} \tag{3}$$

where r_t^e is the instantaneous expected asset return, $r_t^e = \frac{1}{dt}E_{t^-}\left(\frac{dS_t}{S_{t^-}} + \frac{D_t}{S_t}\right)$, D_t is the asset payout, and r_t^f is the instantaneous risk-free rate. In our model, all sources of risk induce a negative co-movement between the asset price and the state-price density. Hence, the asset risk premia load positively on Brownian motion variance, and the good and bad jump intensities: $\Lambda, \beta_g^{ep}, \beta_p^{ep} > 0$.

Next we derive the risk compensation for the variance risks. The total variance of returns, given by the instantaneous quadratic variation in the asset price, is given by

$$V_{t,t+dt} \equiv d[\log S, \log S]_t = \sigma^2\,dt + Z_{g,t}^2\,dN_{g,t} + Z_{b,t}^2\,dN_{b,t}. \tag{4}$$

[7] Section 4 and Online Appendix discuss how the market prices of risk can be linked to preferences and the fundamental cash flow dynamics in a general equilibrium environment.

The quadratic variation consists of a term that represents the diffusive variation, and terms that represent good and bad quadratic jump variation.[8] Notably, the total variance depends on both good and bad jumps: in expectation, it increases in both good or bad jump intensities.

To separately identify the contributions of good and bad jumps, we use the results in Barndorff-Nielsen, Kinnebrock, and Shephard (2010) for the asymptotic behavior of the upside and downside realized variances. The two variances decompose the total variation into two components associated with positive and negative moves in prices, which we describe in a further detail in the empirical part of the paper. In our theoretical setting, the two instantaneous variations are given by

$$V^g_{t,t+dt} = \frac{1}{2}\sigma^2\, dt + Z^2_{g,t}\, dN_{g,t}, \quad V^b_{t,t+dt} = \frac{1}{2}\sigma^2\, dt + Z^2_{b,t}\, dN_{b,t}. \tag{5}$$

The good and bad return variation share a common component due to the symmetric Brownian motion shock, and isolate respective intensities of good and bad asset jumps, respectively.

The compensation for the total variance risks is given by the variance premium, which is the difference between the physical and risk-neutral expectations of $V_{t,t+dt}$:

$$VP_{t,t+dt} = \frac{1}{dt}\left(\mathbb{E}_{t-}\left[V_{t,t+dt}\right] - \mathbb{E}^Q_{t-}\left[V_{t,t+dt}\right]\right) = \beta^{vp}_g\lambda_g + \beta^{vp}_b\lambda_b. \tag{6}$$

Analogous to the decomposition of the total variance, we can decompose the variance premia into the good and bad components:

$$VP^i_{t,t+dt} = \frac{1}{dt}\left(\mathbb{E}_{t-}\left[V^i_{t,t+dt}\right] - \mathbb{E}^Q_{t-}\left[V^i_{t,t+dt}\right]\right) = \beta^{vp}_i\lambda_i. \tag{7}$$

where $i = g, b$. Notably, the variance premia capture the risk premia due to jump components only: the instantaneous variance of Brownian motion is the same under the physical and risk-neutral measure. Further, the decomposition shows that good variance premium separately identifies the

[8]Because the model is *iid*, the variation over a discrete time horizon is proportional to the instantaneous variation.

good jump risk component, while bad variance premium captures the bad jump risk component. Total variance premium, or any other higher-order moments of returns, are functions of both of these components. This motivates using the good and bad variance premia as economically justified, clean measures of the relevant jump risks in the financial markets.

Our model gives unambiguous predictions for the relation between the equity premium and the components of the variance premium. Recall that the equity premium is increasing in both jump intensities. On the other hand, following our discussion of the risk-neutral measure, the bad variance premium loads negatively on bad jump intensity ($\beta_b^{vp} < 0$), while the good variance premium loads positively on good jump intensity ($\beta_g^{vp} > 0$). Indeed, investors dislike bad jumps, and are willing to pay a premium for the exposure to bad jump shocks, so bad variance premium is negative and decreasing in the amount of bad jump risks. On the contrary, good jumps happen in good economic states of high prices, so the good variance premium is positive, and increasing in the amount of good jump risk.

Comparative statics imply that good variance premium is related to asset risk premia with a positive sign, while bad variance premium is negatively associated with the asset risk premia. In the next section, we evaluate the model predictions empirically, and show that they constitute a robust feature of the data.

Finally, our simple model also highlights the importance of considering the components of the variance premia to study the risk and return relationship, as opposed to total variance premium, or the return volatilities themselves. Indeed, the total variance premia conflates the two opposite effects of good and bad jump intensities (see equation (6)), and thus masks the relation between the variance and the risk premium. The return volatilities – good, bad, and total – further depend on the volatility of the Brownian motion, as shown in equation (5). The relationship between the return volatilities and the risk premia are governed by the relative risk and return tradeoff for the good jump, bad jump, and Brownian motion risk. This makes the signs of the slope coefficients in projections of future returns on good and bad volatilities, unlike the good and bad variance premia, ambiguous.

3 Empirical analysis

In this section, we develop measures of the good and bad variance premia and provide empirical evidence for their relation to expected returns. Below, we describe the construction of the benchmark variance measures and provide a discussion of our findings, followed by various robustness checks.

3.1 Good and bad implied variances

Our goal is to construct intuitive measures that separately characterize the variation in the positive and negative components of the conditional aggregate equity return distribution. For this purpose, we utilize the method developed by Bakshi and Madan (2000) and Bakshi, Kapadia, and Madan (2003) to infer the moments of the risk-neutral return distribution from the cross-section of option prices in a model-free way.

We consider the price of a volatility contract that pays off the squared log return at time $t+1$. Let s_t denote the natural logarithm of the price, S_t, of the underlying market index at time t. The payoff of the contract is $r_{t+1}^2 \equiv (s_{t+1} - s_t)^2$. We define total implied variance, iv_t, as the price of the contract:

$$iv_t \equiv e^{-r_t^f} \, \mathbb{E}_t^Q \left[r_{t+1}^2 \right], \tag{8}$$

where $\mathbb{E}_t^Q$ denotes the expectation under the risk-neutral measure conditional on time-t information, and r_t^f is the risk-free rate. Bakshi, Kapadia, and Madan (2003) show that iv_t can be calculated using prices of OTM call and put options:

$$iv_t = \int_{S_t}^{\infty} \frac{2(1 - \log(K/S_t))}{K^2} \, C(t, t+1, K) \, dK + \int_0^{S_t} \frac{2(1 + \log(S_t/K))}{K^2} \, P(t, t+1, K) \, dK, \tag{9}$$

where $C(t, t+1, K)$ and $P(t, t+1, K)$ denote the time-t prices of call and put contracts with a time-to-maturity of one period and a strike price of K.

We can write the payoff of the volatility contract as a sum of the two components associated with the positive and negative returns, respectively,[9]

$$r_{t+1}^2 = r_{t+1}^2 \, \mathbb{I}(r_{t+1} > 0) + r_{t+1}^2 \, \mathbb{I}(r_{t+1} \leq 0). \tag{10}$$

Following this decomposition, we define the good and bad implied variances, iv_t^g and iv_t^b, as the prices of the positive and negative payoff components, respectively,

$$\begin{aligned} iv_t^g &\equiv e^{-r_t^f} \, \mathbb{E}_t^Q \left[r_{t+1}^2 \, \mathbb{I}(r_{t+1} > 0) \right], \\ iv_t^b &\equiv e^{-r_t^f} \, \mathbb{E}_t^Q \left[r_{t+1}^2 \, \mathbb{I}(r_{t+1} \leq 0) \right]. \end{aligned} \tag{11}$$

Because the good and bad components of the payoff add to the total payoff of the volatility contract, we have $iv_t = iv_t^g + iv_t^b$ by no-arbitrage.

Similar to the total volatility contract, the prices of its good and bad components can also be computed in a model-free way from the cross-section of option prices. Indeed, Appendix A.1 shows that iv_t^g and iv_t^b correspond to the first and second integrals in (9), respectively,

$$\begin{aligned} iv_t^g &= \int_{S_t}^{\infty} \frac{2(1 - \log(K/S_t))}{K^2} \, C(t, t+1, K) \, dK, \\ iv_t^b &= \int_0^{S_t} \frac{2(1 + \log(S_t/K))}{K^2} \, P(t, t+1, K) \, dK. \end{aligned} \tag{12}$$

This result is intuitive. Total implied variance iv_t is the weighted sum of the option prices, and the variance components are identified by claims that have payoffs contingent on the sign and magnitude of the realized return. Good implied variance is identified by call options that pay off only in case the return realization is positive, and bad implied variance is characterized by put options that pay off only if a negative return is realized.

We compute the implied variance measures at a monthly frequency using the S&P 500 index options data from OptionMetrics from January 1996 to August 2014. We use the averages of the bid and ask quotes for each option contract and eliminate the options with a mid quote lower

[9] One can consider thresholds different than zero for the decomposition. We opt for an intuitive measure of good and bad events in the stock market and characterize them by positive and negative returns.

than \$ 3/8. Options with zero trading volume, with a time-to-maturity lower than 7 days, and those that violate standard no-arbitrage conditions are also filtered out. For each option contract, we compute moneyness as the strike price divided by the current value of the S&P 500 index, K/S_t. As (12) indicates, we use only OTM options, namely, calls with $K/S_t > 1$ and puts with $K/S_t < 1$. The range of available moneyness in the data highly varies over time. We use options in the moneyness range from 0.85 to 1.15 in our benchmark analysis. This ensures that our implied variance measures capture similar economic content throughout the time series and are not driven by deep OTM options at moneyness levels that are not available most of the time. The integrals in (12) are calculated using a cubic spline across implied volatilities because a continuum of strike prices is not available. If the lowest (highest) available moneyness is higher than 0.85 (lower than 1.15), we use the implied volatility of the lowest (highest) moneyness level for implied volatilities outside the available range in the data. We then compute the option prices from the implied volatilities on the grid of spline. We linearly interpolate across available maturities to compute one-month implied variances. We report the robustness of our results to various alternative choices of the moneyness range in Section 3.5.

Figure 1 plots the time series of total, good, and bad implied variances.[10] The implied variance measures are volatile and spike in bad economic times, such as the Russian financial crisis in 1998, leading to the failure of LTCM, the burst of the dot-com bubble in 2002, and the Great Recession period in 2008 and 2009. Total implied variance behaves very similarly to the traditional measure of return volatility, namely, VIX. Indeed, the correlation between iv_t and VIX^2 is 99.24% in levels, and 97.33% in first differences. Table 1 provides summary statistics for the variance measures, and Table 2 provides a correlation matrix. Both tables report statistics for the sample from January 1996 to August 2014, and a subsample excluding the Great Recession from December 2007 to June 2009. iv_t^b is higher than iv_t^g, both on average and throughout the sample. Good and bad implied variances are highly correlated with a correlation coefficient of 0.97 and are both fairly persistent with AR(1) coefficients of 0.81.

[10]Following Bollerslev, Tauchen, and Zhou (2009), Drechsler and Yaron (2011), and Zhou (2009), we report variances in units of volatility in percent, squared, and divided by 12 for comparability with the existing literature. As a result, all variance measures have the same units as $VIX^2/12$.

3.2 Good and bad realized variances

In addition to the risk-neutral variance measures, we construct measures of the realized variances using high-frequency data. Following Barndorff-Nielsen and Shephard (2004), realized variance, rv_t, is defined as the sum of squared high-frequency returns, $r_{t,i}$, within each period:

$$rv_t = \sum_{i=1}^{H_t} r_{t,i}^2, \tag{13}$$

where H_t is the number of high-frequency intervals in period t.

Following Barndorff-Nielsen, Kinnebrock, and Shephard (2010), we decompose the realized variance, rv_t, into good and bad realized variances, rv_t^g and rv_t^b, as

$$rv_t^g = \sum_{i=1}^{H_t} r_{t,i}^2 \, \mathbb{I}(r_{t,i} > 0), \;\; rv_t^b = \sum_{i=1}^{H_t} r_{t,i}^2 \, \mathbb{I}(r_{t,i} \leq 0). \tag{14}$$

Intuitively, the good and bad realized variance measures capture information about time variation in the positive and negative components of the physical distributions of returns. Indeed, Barndorff-Nielsen, Kinnebrock, and Shephard (2010) show that, as $H_t \to \infty$, rv_t^g and rv_t^b converge to half of the Gaussian diffusion in the returns and positive and negative quadratic jump variation, respectively.

We use 5-minute S&P 500 futures returns from TICKDATA to construct the realized variance measures at a monthly frequency.[11] As shown in Table 1, the total, good, and bad realized variance measures have similar statistical properties and are more volatile and less persistent than their implied counterparts. Figure 2 plots the time series of total, good, and bad realized variances. All of the realized variances move closely with each other, consistent with the correlation evidence in Table 2. The realized variances spike in bad economic times, especially in October 2008, following the Lehman crash. This corresponds to the month with the lowest aggregate equity return in our sample. In terms of the difference between good and bad realized variances, bad variance tends

[11] Liu, Patton, and Sheppard (2012) show that 5-minute sampling frequency achieves the optimal trade-off between the precision of the estimators and the impact of the microstructure noise. We treat returns from the close of a trading day to the open of the next trading day as a 5-minute interval.

to rise above good variance in several periods of variance spikes, such as the Great Recession. Another large deviation of rv_t^b from rv_t^g occurs in September 2001. This is due to one data point: stock exchanges were closed from September 10, 2001 to September 17, 2001. At the opening of the market on September 17, the S&P 500 futures price fell by 5.5%, leading to a large increase in bad, relative to good, realized variance.

3.3 The good and bad variance premia

The variance premium is defined as the difference between the physical and risk-neutral expectation of quadratic return variation. Using the measures of implied and realized variance above, we define our proxy for the variance premium vp_t as the difference between rv_t and iv_t:[12]

$$vp_t = rv_t - iv_t, \tag{15}$$

and the good and bad variance premia, vp_t^g and vp_t^b, as the difference between corresponding realized and implied variances

$$vp_t^g = rv_t^g - iv_t^g, \;\; vp_t^b = rv_t^b - iv_t^b. \tag{16}$$

Figure 3 plots the time series of the total, good, and bad variance premia, and Table 1 reports summary statistics. Total implied variance is higher than realized variance most of the time, so that vp_t is negative on average. Similarly, the bad variance premium is also negative throughout most of the sample. Notably, the absolute value of vp_t^b is larger than that of vp_t on average. On the other hand, good implied variance tends to be lower than good realized variance, so that that our estimate of the average good variance premium is positive.

[12] The empirical measures for the variance premium are only proxies for the actual variance risk premia, as they rely on the statistical estimates of the conditional variance. Our definition of the variance premium corresponds to the one in Bollerslev, Tauchen, and Zhou (2009), who assume $rv_t \approx \mathbb{E}_t[rv_{t+1}]$. For robustness, we also entertain forward-looking measures of expected quadratic variation as discussed in Section 3.5. Alternatively, Bakshi and Kapadia (2003) identify the variance premium from delta-hedged option gains, and Carr and Wu (2009) develop a model-free empirical strategy for identification of the variance premium as the price minus the payoff of a contract that pays off realized variance.

Table 2 shows that in the benchmark sample, the correlation between the good and bad variance premia is 87%. A significant portion of this correlation, however, is driven by a sharp increase in October 2008, caused by a large jump in all the realized variances in this month. Removing this data point, the correlation drops to 59%. Notably, the correlation between the good and bad variance premium measures is significantly smaller than those between the good and bad variances themselves. Indeed, the correlations between the good and bad realized or implied variances exceed 97% in the full sample and are above 93%, excluding the Great Recession. Thus, using the variance premium rather than variance measures helps isolate separate variations in the negative and positive components of the distribution.

The signs and magnitudes of the good and bad variance premia suggest that investors are willing to pay a positive premium for an asset that pays off when bad variance is high. The magnitude of this premium is quantitatively large and higher than the premium for an asset that pays off when total variance is high. On other hand, the average risk compensation to hedge good variance risk is negative. The distinct properties of vp_t^g and vp_t^b suggest that they can potentially contain more information and help identify separate risk factors compared to vp_t alone.

3.4 Predicting returns with good and bad variance premia

In this section, we analyze the predictability of future returns by the variance premium measures. As shown in the literature (see Bollerslev, Tauchen, and Zhou (2009)), the total variance premium has a significant predictive power for aggregate equity market returns at short horizons of 3 to 6 months. We show that good and bad components of the variance premium, vp_t^g and vp_t^b, jointly predict excess returns with statistically significant coefficients at horizons longer than 6 months, up to 24 months with increasing R^2 values in horizon. In multivariate predictability regressions, the coefficient on vp_t^g is positive, while it is negative and larger than that of vp_t for vp_t^b in absolute value. This evidence also holds in corporate bond markets and a cross-section of equity portfolios and is robust to alternative specifications of the benchmark predictability regressions.

All predictability regressions presented in this section have the following form:

$$\frac{12}{h}\sum_{i=1}^{h}(r_{t+i} - r^f_{t+i-1}) = \beta_{0,h} + \beta_h' X_t + \epsilon_{t+h}, \tag{17}$$

where h is the horizon of the regression, r_{t+1} denotes a monthly log return, and r^f_t is the monthly log risk-free rate. The vector X_t contains the predictor variables.

To evaluate the statistical significance of our predictors, we report two types of standard errors and the associated p-values.[13] The first are the Newey-West standard errors with optimal number of lags, as in Newey and West (1994). However, Newey-West standard errors are known to be biased downward in finite samples with overlapping returns (Hodrick (1992)). Another concern is that both the returns and variance premia deviate substantially from the average values during periods of high economic stress, such as during the Great Recession. Because both dependent and independent variables take extreme values, in small samples, this can make standard asymptotic inference unreliable. To alleviate these concerns, we adopt a bootstrap approach that relies on the empirical distribution of returns to compute standard errors. The series are simulated under the null hypothesis of no predictability, preserving the heteroskedasticity of the original return sample in the data, and the same relation between extreme values of predictors and volatility of future returns. This allows us to characterize the small-sample distribution of predictive coefficients and compute corresponding standard errors. Appendix A.2 provides details about the bootstrap approach. Because the small sample distributions are not Normal, we directly report the p-values, rather than t-statistics, for the hypothesis tests.

Our economic model implies that the good variance premium should predict aggregate equity returns with a positive sign, while the coefficient on the bad variance premium is negative. To assess the plausibility of these economic restrictions in the data, we report the p-values for the one-sided tests. Specifically, the p-value for the good (bad) variance premium coefficient is computed under the null hypothesis of no return predictability against an alternative hypothesis that the slope coefficient is positive (negative). For consistency, we also rely on the one-sided p-values to

[13] We consider alternative approaches to compute standard errors in Section 3.5.

test return predictability by the total variance premium. For equity returns, we test whether the coefficient on vp_t is negative, consistent with the literature and our economic model. For other assets, such as corporate bonds, the signs of the coefficients do not follow directly from the theory. In this case, we follow a rule that if a predictive coefficient is positive (negative) at the 12-month horizon in our benchmark analysis in Table 3, the p-value corresponds to the probability of a higher (lower) value for the coefficient under the null of no predictability.

Predictability evidence

Panel A of Table 3 reports the results of predictive regressions for excess aggregate equity returns by vp_t, as well as vp_t^g and vp_t^b, jointly. We find that our proxy for the total variance premium, vp_t, predicts returns with a negative sign, and the predictive coefficient is statistically significant at 1- and 3-month horizons.[14] The goodness of fit peaks at a 3-month horizon with an R^2 of 6%.[15] The predictive coefficient is not statistically different from zero for the horizons of 6 months and above under the bootstrap standard errors and for 12 months and above using the Newey-West standard errors. At these horizons, the R^2 values are essentially zero.[16]

Next, we turn to the multivariate predictability regressions using vp_t^g and vp_t^b. At short horizons, the predictive power of the two variance premia is comparable to that of vp_t alone. However, at longer horizons, the evidence suggests that the two variance premia contain additional, economically and statistically significant information about the expected returns. First, the coefficient on vp_t^g is positive at all horizons, starting from 3 months. It is statistically significant under the Newey-West standard errors from 6 months, and under the bootstrap standard errors from 12 months. The coefficient on the bad variance premium is always negative, and is significant for the horizons of 3 months and above. Notably, the absolute value of the coefficient on the bad variance premium in the multivariate setting is significantly larger than that on the total variance

[14]We refer to coefficients with bootstrap p-values smaller than or equal to 5% as statistically significant unless stated otherwise.

[15]All reported R^2 values are adjusted.

[16]Our findings are comparable to those in the literature that rely on VIX, rather than implied variance, to compute variance premium (see Bollerslev, Tauchen, and Zhou (2009)). In this case, the coefficient on total variance premium is significant up to 12 months, and the adjusted R^2 values at 1-, 3-, 6-, 12-, 18-, and 24-month horizons are 5%, 8%, 4%, 1%, 0%, and 0%, respectively.

premium in univariate regressions: it is about twice as large, even at a 1-month horizon. Finally, the multivariate regression delivers a significantly higher goodness of fit, with R^2 values of 7% at the 12-month and 9% at the 24-month horizons, while there is no evidence of significant predictive power by vp_t alone at these horizons.

Panels B and C of Table 3 show predictability evidence for high-yield and investment-grade bond returns from Barclay's bond indexes. Corporate bonds are risky securities that are claims on firms' assets, so we use them as alternative test portfolios to expand the predictability evidence. At horizons of 1 and 2 years, the total variance premium can predict future excess returns with an R^2 of 4% for high-yield portfolios and 6% for investment-grade portfolios. Interestingly, the signs of the coefficients on vp_t are positive at these horizons and statistically significant. Similar to aggregate equity results, decomposing the total variance premium into its good and bad components substantially magnifies bond return predictability. At the horizons of 1 and 2 years, the R^2s in multivariate regressions range from 18% to 24% for high-yield and investment-grade bond returns. The coefficients on the good variance premium are all positive, and those on the bad variance premium are negative. Nearly all of the coefficients are significant under both the Newey-West and bootstrap standard errors. Notably, compared with equity returns, the coefficients on the good variance premium are often larger, in absolute value, than those on the bad variance premium, which can help explain the positive coefficient on the total variance premium in univariate regressions.

Figure 4 plots the predicted 12-month excess equity returns, implied by our regressions. The predicted excess equity returns implied by the multivariate regression stay at low levels in moderate periods, such as from 2003 to 2007, and spike in periods of distress, such as during the Great Recession period in 2008 and 2009. For comparison, we also show the predicted 12-month excess returns from the univariate regression based on the total variance premium. The implied risk premium is much less volatile and does not capture pronounced increases in risk in periods of high economic stress. The correlation between the implied equity premium extracted from univariate and multivariate projections is quite low, and is equal to 15%. We also verify that the results are similar for bond returns. The predicted bond returns implied by the univariate regression are less volatile than those based on the multivariate specification, and the correlations between the two are quite low.

We expand the set of the test assets to include the cross-section of equity portfolio returns. This allows us not only to assess the predictability evidence at the level of individual portfolios but also to run a *joint* test that predictability coefficients across *all* of the portfolios on the good (bad) variance premium are positive (negative).[17] Table 4 summarizes the results for ten size, book-to-market, and industry portfolios. The cross-sectional results support and strengthen our evidence for the aggregate returns. At a 3-month horizon, the total variance premium is a significant predictor in the cross-section of book-to-market portfolios, with a joint p-value of 1%, and is marginally significant for size and industry portfolios with p-values of 8%. The median R^2s vary from 3% for industry to 6% for book-to-market portfolios. There is, however, no significant evidence of predictability at long horizons: the p-values are above 40% at the 1-year horizon and 60% at the 2-year horizon, and all the R^2 are below 1%. On the other hand, the coefficients on the good and bad variance premia are statistically significant, in joint tests, at these horizons, with p-values of 4% and below. The median R^2 increases from 8% at a 1-year to 19% at a 2-year horizon for size portfolios, from 3% to 9% for book-to-market, and from 5% to 8% for industry portfolios.

Our evidence suggests that the variance premium is driven by at least two factors that have opposite association with long-horizon excess returns. This can account for the weaker predictive power of vp_t alone compared to vp_t^g and vp_t^b, jointly. Once the two factors of the variance premium are uncovered, predictability becomes stronger at longer horizons, suggesting that the variance premium contains information about persistent risk factors in financial markets. Further, both the good and bad variance premia need to be included to capture the variation in bond and equity risk premia. Using either of the two in the univariate regressions considerably lowers the amount of return predictability. For example, at 1 year horizon, bad variance premium by itself predicts future stock, investment bond, and high-yield bond returns with R^2s of 1%, 0%, and 2%,

[17] Appendix A.2 describes the computation of p-values for this exercise.

respectively; the corresponding R^2s for the good variance premia are 0%, 11%, and 9%.[18]

Finally, our benchmark empirical evidence is related to the recent literature which highlights the link between the higher-order moments of returns and the asset prices. As we discussed in Section 2, in our illustrative model all the higher-order moments of equity prices are spanned by the two state variables, and thus, by the two variance premia. However, the two variance premia are able to separately identify the movements in the primitive uncertainties about the left and right tails of the return distribution. Higher-order moments, on the other hand, confound the two state variables, and thus do not provide a clean interpretation of the economic and statistical roles of these uncertainty channels. In a related setup, Feunou, Jahan-Parvar, and Okou (2014) note that the coefficients on the two variance premia are similar, in absolute values, and use the difference between the two premia as a measure of the skewness risk premium. Using our estimates, we can reject the null hypothesis that the two slope coefficients sum up to zero for all the assets and nearly all the horizons using the Newey-West standard errors. Under the bootstrap standard errors, we can reject it for investment bond and high-yield bond returns. For equity returns, the null hypothesis that the two slope coefficients are the same, in absolute value, is rejected in the sample which excludes the period of a large market turmoil around the crisis period.

Further, we directly construct the implied skewness measure, *skew* following Bakshi, Kapadia, and Madan (2003) and also compute realized skewness using the high-frequency return data as in Amaya, Christoffersen, Jacobs, and Vasquez (2013). The difference between the realized and implied skewness is a measure of the skewness risk premium, *srp*.[19] The Online Appendix shows that adding either *skew* or *srp* to the good and bad variance premium measures does not affect our benchmark evidence for the relation between the variance premium and the future returns. In fact, the significance of the coefficients on the variance premia tends to improve when we control for measures of skewness. The coefficients on skewness tend to be negative for equity and positive

[18] Expanding the range of options to construct the implied variances helps increase the predictive R^2s for the bad variance premia, consistent with findings in Bollerslev, Todorov, and Xu (2015) that very deep out-of-the-money put options contain information about equity premium.

[19] Chang, Christoffersen, and Jacobs (2013) show evidence of a skewness risk premium in the cross-section of returns, and Conrad, Dittmar, and Ghysels (2013) find that individual stocks' skewness is strongly related to future returns. Kozhan, Neuberger, and Schneider (2013) find that a skewness premium can explain most of the implied volatility skew for index options.

for bonds, but are quite imprecisely estimated.

In all, our empirical results suggest that both uncertainties about good and bad market events capture sizeable and separate information about the variations in asset risk premia.

3.5 Robustness tests

In this section we show that our empirical evidence is robust to different sample periods, truncation of the outliers, alternative specifications for the predictive regressions, and construction of the variance premium measures.

Alternative measurements of variances

We check the robustness of our results to a number of alternative measurements of the variances and the variance premia.

Our realized variance measures are constructed in a standard way using the variation in high-frequency 5-minute returns, while the implied variances are based on the risk-neutral expectations of next month's squared return. To make sure that our results are not driven by the mismatch in the sampling frequencies underlying the two variance measures, we consider alternative measures of physical variances which are based on the ex-ante variation in monthly returns. Specifically, we consider linear projections of next-month squared returns, multiplied by the corresponding sign indicators, on the available good and bad, realized and implied, variance measures:

$$\begin{aligned} r_{t+1}^2 &= \alpha_0 + \alpha_1 \left[iv_t^g \quad iv_t^b \quad rv_t^g \quad rv_t^b \right]' + \epsilon_{t+1}, \\ r_{t+1}^2 \mathbb{I}(r_{t+1} > 0) &= \alpha_0^g + {\alpha_1}^g \left[iv_t^g \quad iv_t^b \quad rv_t^g \quad rv_t^b \right]' + \epsilon_{t+1}^g, \\ r_{t+1}^2 \mathbb{I}(r_{t+1} \leq 0) &= \alpha_0^b + {\alpha_1}^b \left[iv_t^g \quad iv_t^b \quad rv_t^g \quad rv_t^b \right]' + \epsilon_{t+1}^b. \end{aligned} \tag{18}$$

We use the fitted values from these regressions to construct ex-ante total, bad, and good physical variance measures, respectively, which in turn are used to define alternative measures of total, bad, and good variance premia. In Table 5 we show that the predictive evidence based on these alternative measures of the variance premia is very similar (in fact, in most cases even stronger) than in the benchmark specification.

We consider various other variations in the construction of the variance measures, and report our findings in the Online Appendix. For implied variances, we extend the range of the cubic spline to 0 - 3, use options in the moneyness range 0.8 - 1.2 and 0.9 - 1.1. We also compute realized variances using 5-minute returns of S&P 500 E-Mini futures and use expected, rather than realized, variances to construct the variance premium measures. Overall, our key results remain robust to these alternative specifications.

Alternative standard errors

In our benchmark analysis we report two types of standard errors and the associated *p*-values. One is the Newey-West standard errors with optimal number of lags based on Newey and West (1994), and the other is the bootstrap approach which preserves the heteroskedasticity of returns and the relation between extreme values of predictors and the volatility of returns. We also consider Hodrick (1992) 1B standard errors, and standard errors using the ARMA(1,1)-GARCH(1,1) parametric bootstrap approach. As pointed out in Bakshi, Panayotov, and Skoulakis (2011), this parametric bootstrap approach preserves the negative MA component in variance premia.

As shown in Table 6, these alternative approaches deliver similar results to our benchmark analysis. The variance premium is insignificant after 3 to 6 months, while the coefficients on good and bad variance premia are significant beyond 6 months in all predictive regressions using equity and bond returns. At these horizons, Hodrick (1992) standard errors are generally larger than the Newey-West, but the inference is not as conservative compared to our benchmark bootstrap approach. The ARMA(1,1)-GARCH(1,1) parametric bootstrap of Bakshi, Panayotov, and Skoulakis (2011) generally leads to similar (and often stronger) conclusions for the predictive power of the good and bad variance premia relative to the non-parametric bootstrap.

Out of sample predictions

A common concern with predictability regressions is that they do not work well out of sample (see Goyal and Welch (2008)). We test the robustness of our results out of sample using the out-of-sample R^2. The out-of-sample R^2 is calculated by comparing the mean-squared errors for predicting future returns using the predictive model versus the average historic estimate of the premium. For robust predictors, the out-of-sample R^2s should be greater than zero and similar to the in-sample counterparts. The formal significance of the out-of-sample R^2 can be assessed by the p-value test of Clark and West (2007).

The out-of-sample evidence is consistent with our benchmark story. For aggregate equity, using the total variance premium leads to out-of-sample $R^2 s$ of 7-8% up to a 3-month horizon based on 120-month rolling estimates. The out-of-sample evidence deteriorates at longer horizons, with negative R^2s and p-values of above 30%. On the other hand, using good and bad variance premia produces out-of-sample R^2s of 7% and 24%, respectively, at 1 and 2 year horizons, and the p-values are below 1%. The results are even more pronounced for bond returns. The out-of-sample evidence is insignificant at all horizons using single variance premium predictor. Using good and bad variance premia produces out-of-sample R^2s of 20-30%, and p-values below 3% at horizons of 3 months and beyond.

Long sample

Our benchmark sample covers the period from January 1996 to August 2014, given the availability of options data from OptionMetrics. To provide further robustness checks to our main results, we extend the sample using options data from CBOE from January 1988 to December 1995.[20] The options in this data set are not as liquid as those in the OptionMetrics sample from January 1996, and the range of available moneyness is significantly narrower. Because of these data issues, we choose to use these data in the robustness checks and focus on a more standard data set from OptionMetrics in the benchmark analysis.

As shown in Table 7, the empirical results based on the extended sample are quite similar to the benchmark sample. In univariate regressions, the total variance premium is a significant

[20]This data set is also used in Ait-Sahalia, Wang, and Yared (2001) and Eraker (2004), among others. We thank Bjorn Eraker for providing this data set to us. We follow the OptionMetrics procedure to compute dividend yields and the zero-coupon yield curve for this data set. We follow the OptionMetrics procedure for dividend yields and risk-free rate used to compute implied volatilities from January 1988 to December 1995.

and predictor of future equity returns with a negative sign from 1 to 6 months, while we do not find any evidence for predictability of future bond returns. The R^2s do not exceed 5% in equity return regressions. In multivariate regressions, a high good (bad) variance premium predicts higher (lower) future asset returns across all of the maturities above 1 month. The significance of the coefficients is comparable to that in the full sample. The R^2s are 4%, 14% and 11% for stock, high-yield and investment-grade bond returns, respectively, at the 1-year horizon, all of which are substantially above those for the total variance premium alone.

Notably, our main findings are not driven solely by the benchmark sample, which starts in 1996, and they hold even in the restricted sample from 1988 to 1995. Indeed, in this subsample, the 1-year predictive coefficients on the good and bad variance premia are 0.25 and -0.42, respectively, so that the good (bad) variance premium predicts future equity returns with a positive (negative) sign. Results are similar for the bond returns.

Robustness to outliers

The Great Recession (December 2007 - June 2009) was a period of high equity market volatility, as shown in Figures 1 and 2, which translates into the extreme values for the variance premium measures (see Figure 3). We perform several robustness checks to assess whether these observations play a significant role in our results.

First, we exclude the Great Recession period from our analysis. The predictability results for the sample excluding December 2007 - June 2009 are reported in Table 8. For aggregate equity returns, the results are similar to those based on the entire sample, both in terms of the magnitude and sign of coefficients and an increase in predictive power in multivariate regressions. In this restricted sample, vp_t predicts returns with a negative coefficient, which is statistically significant at all horizons as opposed to the full sample, where the predictive coefficients are not significant at long horizons. However, the goodness of fit tends to be quite low, and the R^2s do not exceed 2%. In multivariate regressions, the signs on the good (bad) variance premium are all positive (negative), in line with the full sample evidence. While the coefficient on vp_t^b is significant at all horizons above 3 months, the significance of the coefficient on vp_t^g is weaker compared to the full sample regressions. The bootstrap p-values are 8% and 5% at 12- and 24-month horizons, respectively, whereas the Newey-West p-values are below 4%. For corporate bonds, we find no evidence of return predictability by vp_t at any horizon, and the R^2 values are essentially zero. At the same time, the joint predictive power of vp_t^g and vp_t^b increases with horizon, and signs are in line with the full sample. At 1- and 2-year horizons, the R^2s are 6% and 7% for high-yield bonds, and 2% and 5% for investment-grade bonds, respectively. The coefficients are significant for high-yield bonds at most horizons, while they are significant for investment-grade bonds at a 3-month horizon, and marginally significant at a 2-year horizon.

We also test the robustness of the predictability results for the cross-section of equity portfolios to the exclusion of the Great Recession period and find similar results. For all three groups of portfolios considered in Section 3.4, the predictive coefficient on the total variance premium is significant at 3-, 12-, and 24-month horizons, however, with moderate R^2 values that do not exceed 3% at any of the horizons considered for any of the portfolios. In multivariate regressions, the coefficients on the good (bad) variance premium are significantly positive (negative) with p-values below 2% at 3-, 12-, and 24-month horizons. The predictive power in multivariate regressions is stronger, in line with the full sample evidence.

Another way to assess the importance of the outliers is to truncate the positive (negative) observations for the bad (good) variance premia. Indeed, say, large positive values for the bad variance premium during the crisis most likely reflect measurement issues in capturing the ex-ante variances by the realized variances, and are hard to rationalize in a risk-based economic model that we consider. To assess the importance of these observations, we check the benchmark predictive evidence using the truncated variance premia: $vp^{g,*} = max(vp^g, 0)$, and $vp^{b,*} = min(vp^b, 0)$. The results are very similar to the benchmark setting: for example, at 1 year, the predictive R^2s for stock, investment-grade bond, and high-yield bond returns are 6%, 16%, and 17%.

Finally, we also verify the robustness of our results to outliers by running a robust predictability regression. In this approach, we exclude any observation with a Cook's distance greater than one and use Huber (1964) weights to determine the coefficients. The results are consistent with the benchmark findings and are omitted for brevity.

Constrained predictability regressions

Our benchmark estimate of the implied equity premium, obtained from a linear regression of future 12-month returns on the two variance premia, occasionally takes negative values, as can be seen from Figure 4. Our economic model, however, suggests that the conditional equity premium should always be positive. To assess the role of the negative estimates for the implied equity premium for the predictability evidence, we follow Pettenuzzo, Timmermann, and Valkanov (2014) and consider a constrained regression approach in which the linear regression coefficients are estimated subject to the constraint that the fitted values are positive at each point in the sample.

Notably, the largest observed negative equity premium corresponds to September 2001. As discussed in Section 3.2, this is due to the stock market interruption for a week following September 11, and a mechanically inflated value for the realized bad variance, which treats the realized weekly return as a 5-minute return. Although this measurement outlier does not affect our benchmark results, it impacts constrained regressions, which are sensitive to extreme observations. Because this observation is primarily caused by statistical issues in measurements of the realized variance, rather than economic considerations, we exclude it from the constrained regression specification.[21]

Table 9 reports the predictability results for the constrained specification, and Figure 4 plots the predicted series from multivariate regressions. In univariate regressions, the constrained approach tries to address the negativity of the implied equity premium in October 2008 by increasing the constant of the regression and increasing the coefficient on vp_t to virtually zero. As a result, in the constrained estimation there is no significant evidence of stock return predictability by total variance risk premium at all horizons. The predicted excess return is also negative in the multivariate setting in this period. However, because the coefficients on vp_t^g and vp_t^b are of the opposite sign, and both variance premium measures spike sharply in this period, the required adjustment on the coefficients is much smaller. Overall, the predicted values from constrained and unconstrained regressions closely track each other as shown in Figure 4. The statistical

[21] Specifically, bad realized variance is 74.37 in September 2001. Almost half of this value, 32.02, is attributable to the week when markets were closed. If we assume that the weekly price drop occurred in 5, 10, or 400 intervals (there are around 400 5-minute observations during the week), the contribution to bad realized variance would be 6.40, 3.20, or 0.08 instead of 32.02. All of these values lead to a positive fitted value for the excess return.

predictability evidence also remains very similar to the benchmark setting, as shown in Table 9.

4 Economic Interpretation

We develop a consumption-based general equilibrium framework that can provide economic intuition for the illustrative model in Section 2. We include a brief summary of the model in this section, and leave details of the model setup and solution to the Online Appendix.

In line with Segal, Shaliastovich, and Yaron (2015), Tsai and Wachter (2015) and Bekaert and Engstrom (2015), we consider good and bad jumps in the consumption dynamics with time-varying jump intensities, which drive fluctuations in the quantities of upside and downside jump risk. Investors have recursive preferences of Epstein and Zin (1989) and Weil (1990) with an elasticity of inter-temporal substitution (EIS) equal to one. When the risk aversion is above one, our equilibrium model delivers the key intuition of the illustrative model that good jumps, both in cash flow levels and jump intensities, represent good states of the economy, while bad jumps are bad economic states. The parameter configuration also ensures that the equilibrium equity prices increase with good jumps and fall with bad jumps, which is again consistent with the reduced-form specification of the illustrative model.

We calibrate the model and show that it provides a good fit to the key moments of the macroeconomic, equity price, and the variance premium data, including the evidence for the predictive power and the informational content of the variance premia for future equity returns. We further use the calibrated model to infer the implied values of the good and bad jump intensities from the time series of the good and bad variance premia in the data, and then compute the model-implied estimate of the conditional equity premium. We plot the model equity premium along with the predicted 12-month excess returns from the constrained regression approach in Figure 5. Quite remarkably, the model tracks the estimate of the equity premium in the data very well. Indeed, the correlation between the model-implied equity premium and the predicted 12-month excess returns is 90%. The good fit is supportive of the primitive model channels, which support the reduced-form structure of the illustrative model.

5 Conclusion

In this paper, we show that the variance premium is driven by "good" and "bad" components, which separately capture the time-varying risk compensations for the realized variation in positive and negative market returns. The good and bad variance premia contain distinct information about future excess returns on equity and corporate bonds, especially at long horizons of 1 and 2 years.

To rationalize the empirical evidence in the data, we consider a model that features upward (good) and downward (bad) jumps in the price of equity with time-varying intensities. We show that when good (bad) jump shocks carry a positive (negative) market price of risk, the equity premium is increasing in good and bad jump intensities, while the variance premium is increasing in good and decreasing in bad jump intensity. Therefore, the good variance premium predicts future returns with a positive sign, while the bad variance premium predicts them with a negative sign, consistent with the data. In the Online Appendix we quantify the magnitudes of these effects by imposing the restrictions of a general equilibrium model.

There are several extensions of our paper that would be interesting to pursue in future work. On the empirical side, we consider the predictability of the U.S. equity and corporate bond returns. It would be useful to extend our analysis for international asset returns, as in Bollerslev, Marrone, Xu, and Zhou (2014), and for other asset classes, such as government bonds, currencies, and commodities. Further, while our variance premium measures capture volatility risks in aggregate equity markets, one can use similar methods to compute the good and bad variance premia in other asset markets.[22] On the model side, our specification can be extended to allow for separate sources of risk in the expected growth or consumption volatility, as in Bansal and Yaron (2004), including Poisson jump risk as in Drechsler and Yaron (2011), or time-varying skew-normal risk as in Colacito, Ghysels, and Meng (2013). One can also enrich the model by separating volatility and jump risk factors, consistent with the findings in Santa-Clara and Yan (2010), and incorporating additional volatility-related factors, such as the long-run volatility (see Duffie, Pan, and Singleton

[22]See Mueller, Vedolin, and Yen (2013), Londono and Zhou (2014) and Prokopczuk and Wese Simen (2014), who compute the total variance premium for bond, currency and commodity markets, respectively.

(2000)), or volatility of volatility (see Bollerslev, Tauchen, and Zhou (2009)). Further, for simplicity, the two jump intensities are assumed to be uncorrelated in our framework. In the data, the good and bad volatility of returns have a more complex dynamic dependency, as discussed in Patton and Sheppard (2013), and the model can be extended to take this into account.

Appendix

A Empirical analysis

A.1 Good and bad implied variances

Let S be the price of the underlying security at time $t+1$ and $H(S)$ a twice differentiable payoff function. Bakshi, Kapadia, and Madan (2003) show that $H(S)$ can be spanned as

$$H(S) = H(\bar{S}) + (P - \bar{S})H_P(\bar{S}) + \int_{\bar{S}}^{\infty} H_{SS}(K)\ (S-K)^+\ dK + \int_0^{\bar{S}} H_{SS}(K)\ (K-S)^+\ dK, \tag{A.1}$$

where subscripts represent derivatives and $\bar{S}$ is a constant. The payoff of the volatility contract is given by $H(S) = \left(\log\left(\frac{S}{S_t}\right)\right)^2$. The definition of the volatility contract payoff along with (A.1) implies

$$H(S) = \int_{S_t}^{\infty} \frac{2(1-\log(K/S_t))}{K^2}\ (S-K)^+\ dK + \int_0^{\bar{S}} \frac{2(1+\log(S_t/K))}{K^2}\ (K-S)^+\ dK. \tag{A.2}$$

Note that the price of the volatility contract is $iv_t = e^{-r_t^f}\mathbb{E}_t^Q\left[H(S)\right]$, while the price of a call option is $C(t,t+1,K) = e^{-r_t^f}\mathbb{E}_t^Q\left[(S-K)^+\right]$ and the put price is $P(t,t+1,K) = e^{-r_t^f}\mathbb{E}_t^Q\left[(K-S)^+\right]$. Equation (A.2) along with these definitions implies (9).

Now consider the "good" volatility contract with price $iv_t^g = e^{-r_t^f}\ \mathbb{E}_t^Q\left[r_{t+1}^2\ \mathbb{I}(r_{t+1} > 0)\right]$. The payoff can be written as $H^g(S) = \left(\log\left(\frac{S}{S_t}\right)\right)^2\ \mathbb{I}(S > S_t)$. Although H_{SS} is not continuous at $S = S_t$, we can write

$$\begin{aligned} H^g(S) &= \lim_{T \downarrow S_t}\int_T^{\infty} H_{SS}(K)\ (S-K)^+\ dK + \lim_{T\uparrow S_t}\int_0^T H_{SS}(K)\ (K-S)^+\ dK, \\ &= \int_{S_t}^{\infty} \frac{2}{S^2}\left(1 - \log\left(\frac{S}{S_t}\right)\right)\ (S-K)^+\ dK. \end{aligned} \tag{A.3}$$

The computation of good implied variance in (12) follows from (A.3) and the definition of a call option. The computation is analogous for implied bad variance.

A.2 Bootstrap approach

To bootstrap standard errors, we first standardize excess returns, $r_{t+1} - r_t^f$, by subtracting the unconditional mean and dividing by iv_t. We draw a random sample with replacement from this distribution with the same length as the return series in the data. Simulated returns then correspond to standardized returns scaled up by the observations of iv_t. We run the return predictability regressions in the simulated samples on the variance premium measures in the data. We repeat this procedure 10,000 times to construct a small-sample distribution of regression coefficients, which we use to compute standard errors and p-values. In this non-parametric approach, the time series of predictor variables, as well as heteroskedasticity of returns, remain the same as in the data, while there is no predictability of returns under the null.

We further consider the joint evidence for return predictability across multiple equity portfolios, namely, ten portfolios formed on size and book-to-market and ten industry portfolios from Kenneth French's website. Following MacKinnon (2009), we compute p-values that summarize the statistical significance of portfolio return predictability within each of the three groups. We standardize all excess portfolio returns by their unconditional mean and iv_t as described above. From the standardized sample, we draw a random sample of ten returns with replacement with the same length as in the data, and we compute simulated returns using the unconditional average portfolio return and the time series of iv_t. This preserves the dependence of portfolio returns in the cross-section and the heteroskedasticity in the time series. For a predictor variable that predicts aggregate equity market returns with a positive (negative) sign, we choose the maximum (minimum) of the ten predictive coefficients as the critical value. Following this approach, we construct 10,000 cross-sections of returns that have the same size as the original data and create a small-sample distribution of minimum (maximum) coefficient values. The p-value is the probability that the minimum (maximum) predictive coefficient is smaller (larger) than its empirical counterpart in simulated samples.

To compute standard errors based on the parametric bootstrap approach, we closely follow Bakshi, Panayotov, and Skoulakis (2011). We impose an ARMA(1,1)-GARCH(1,1) specification on variance premia (vp_t, vp_t^g, and vp_t^b), which accounts for GARCH effects and the negative MA components in variance premia estimates. See Appendix C of Bakshi, Panayotov, and Skoulakis (2011) for further details.

References

Ait-Sahalia, Yacine, Yubo Wang, and Francis Yared, 2001, Do option markets correctly price the probabilities of movement of the underlying asset?, *Journal of Econometrics* 102, 67–110.

Amaya, Diego, Peter Christoffersen, Kris Jacobs, and Aurelio Vasquez, 2013, Does realized skewness predict the cross-section of equity returns?, Working Paper.

Andersen, Torben G, Luca Benzoni, and Jesper Lund, 2002, An empirical investigation of continuous-time equity return models, *The Journal of Finance* 57, 1239–1284.

Ang, Andrew, Joseph Chen, and Yuhang Xing, 2006, Downside risk, *Review of Financial Studies* 19, 1191–1239.

Bakshi, Gurdip, Charles Cao, and Zhiwu Chen, 1997, Empirical performance of alternative option pricing models, *The Journal of Finance* 52, 2003–2049.

Bakshi, Gurdip, and Nikunj Kapadia, 2003, Delta-hedged gains and the negative market volatility risk premium, *Review of Financial Studies* 16, 527–566.

Bakshi, Gurdip, Nikunj Kapadia, and Dilip Madan, 2003, Stock return characteristics, skew laws, and the differential pricing of individual equity options, *Review of Financial Studies* 16, 101–143.

Bakshi, Gurdip, and Dilip Madan, 2000, Spanning and derivative-security valuation, *Journal of Financial Economics* 55, 205–238.

Bakshi, Gurdip, George Panayotov, and Georgios Skoulakis, 2011, Improving the Predictability of Real Economic Activity and Asset Returns with Forward Variances Inferred from Option Portfolios, *Journal of Financial Economics* 100, 475–95.

Bali, Turan G, and Hao Zhou, 2014, Risk, uncertainty, and expected returns, Working paper.

Bansal, Ravi, and Amir Yaron, 2004, Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles, *Journal of Finance* 54, 1481–1509.

Barndorff-Nielsen, Ole E, Silja Kinnebrock, and Neil Shephard, 2010, Measuring downside risk-realised semivariance, Working Paper.

Barndorff-Nielsen, Ole E, and Neil Shephard, 2004, Power and bipower variation with stochastic volatility and jumps, *Journal of Financial Econometrics* 2, 1–37.

Bekaert, Geert, and Eric Engstrom, 2015, Asset return dynamics under bad environment good environment fundamentals, *Journal of Political Economy, Forthcoming.*

Bekaert, Geert, Eric Engstrom, and Andrey Ermolov, 2015, Bad environments, good environments: A non-gaussian asymmetric volatility model, *Journal of Econometrics* 186, 258–275.

Bollerslev, Tim, James Marrone, Lai Xu, and Hao Zhou, 2014, Stock return predictability and variance risk premia: statistical inference and international evidence, *Journal of Financial and Quantitative Analysis* 49, 633–661.

Bollerslev, Tim, George Tauchen, and Hao Zhou, 2009, Expected stock returns and variance risk premia, *Review of Financial Studies* 22, 4463–4492.

Bollerslev, Tim, Viktor Todorov, and Lai Xu, 2015, Tail risk premia and return predictability, *Journal of Financial Economics* 118, 113–134.

Broadie, Mark, Mikhail Chernov, and Michael Johannes, 2007, Model specification and risk premia: Evidence from futures options, *The Journal of Finance* 62, 1453–1490.

Campbell, John Y, and Ludger Hentschel, 1992, No news is good news: An asymmetric model of changing volatility in stock returns, *Journal of Financial Economics* 31, 281–318.

Campbell, John Y, and Robert J Shiller, 1988, The dividend-price ratio and expectations of future dividends and discount factors, *Review of Financial Studies* 1, 195–228.

Carr, Peter, and Liuren Wu, 2009, Variance risk premiums, *Review of Financial Studies* 22, 1311–1341.

Chang, Bo Young, Peter Christoffersen, and Kris Jacobs, 2013, Market skewness risk and the cross section of stock returns, *Journal of Financial Economics* 107, 46–68.

Clark, Todd, and Kenneth West, 2007, Approximately normal tests for equal predictive accuracy in nested models, *Journal of Econometrics* 138, 291–311.

Cochrane, John H, 2008, The dog that did not bark: A defense of return predictability, *Review of Financial Studies* 21, 1533–1575.

Colacito, Riccardo, Eric Ghysels, and Jinghan Meng, 2013, Skewness in expected macro fundamentals and the predictability of equity returns: Evidence and theory, Working Paper.

Conrad, Jennifer, Robert F Dittmar, and Eric Ghysels, 2013, Ex ante skewness and expected stock returns, *The Journal of Finance* 68, 85–124.

Drechsler, Itamar, and Amir Yaron, 2011, What's vol got to do with it, *Review of Financial Studies* 24, 1–45.

Duffie, Darrell, Jun Pan, and Kenneth Singleton, 2000, Transform analysis and asset pricing for affine jump-diffusions, *Econometrica* 68, 1343–1376.

Epstein, Larry, and Stan Zin, 1989, Substitution, risk aversion and the temporal behavior of consumption and asset returns: A theoretical framework, *Econometrica* 57, 937–969.

Eraker, Bjørn, 2004, Do stock prices and volatility jump? Reconciling evidence from spot and option prices, *The Journal of Finance* 59, 1367–1404.

Eraker, Bjørn, 2008, The volatility premium, Working paper.

Eraker, Bjørn, Michael Johannes, and Nicholas Polson, 2003, The impact of jumps in volatility and returns, *The Journal of Finance* 58, 1269–1300.

Fama, Eugene F, and Kenneth R French, 1989, Business conditions and expected returns on stocks and bonds, *Journal of Financial Economics* 25, 23–49.

Feunou, Bruno, Mohammad R Jahan-Parvar, and Cedric Okou, 2014, Downside Variance Risk Premium, Working Paper.

Feunou, Bruno, Mohammad R Jahan-Parvar, and Roméo Tédongap, 2013, Modeling market downside volatility, *Review of Finance* 17, 443–481.

French, Kenneth R, G William Schwert, and Robert F Stambaugh, 1987, Expected stock returns and volatility, *Journal of Financial Economics* 19, 3–29.

Glosten, Lawrence R, Ravi Jagannathan, and David E Runkle, 1993, On the relation between the expected value and the volatility of the nominal excess return on stocks, *The Journal of Finance* 48, 1779–1801.

Goyal, Amit, and Ivo Welch, 2008, A comprehensive look at the empirical performance of equity premium prediction, *Review of Financial Studies* 21, 1455–1508.

Guo, Hui, Kent Wang, and Hao Zhou, 2014, Good Jumps, Bad Jumps, and Conditional Equity Premium, Working paper.

Han, Bing, and Yi Zhou, 2012, Variance risk premium and cross-section of stock returns, Working paper.

Hodrick, Robert J, 1992, Dividend yields and expected stock returns: Alternative procedures for inference and measurement, *Review of Financial Studies* 5, 357–386.

Huber, Peter J, 1964, Robust estimation of a location parameter, *The Annals of Mathematical Statistics* 35, 73–101.

Keim, Donald B, and Robert F Stambaugh, 1986, Predicting returns in the stock and bond markets, *Journal of Financial Economics* 17, 357–390.

Kelly, Bryan, and Hao Jiang, 2014, Tail risk and asset prices, *Review of Financial Studies* 27, 2841–2871.

Kozhan, Roman, Anthony Neuberger, and Paul Schneider, 2013, The skew risk premium in the equity index market, *Review of Financial Studies* 26, 2174–2203.

Lettau, Martin, Matteo Maggiori, and Michael Weber, 2014, Conditional risk premia in currency markets and other asset classes, *Journal of Financial Economics* 114, 197–225.

Liu, Lily Y, Andrew J Patton, and Kevin Sheppard, 2012, Does anything beat 5-minute RV? A comparison of realized measures across multiple asset classes, Working Paper.

Londono, Juan M, and Hao Zhou, 2014, Variance risk premiums and the forward premium puzzle, Working Paper.

MacKinnon, James G, 2009, Bootstrap hypothesis testing, *Handbook of Computational Econometrics* pp. 183–213.

Mueller, Philippe, Andrea Vedolin, and Yu-min Yen, 2013, Bond variance risk premiums, Working paper.

Newey, Whitney K., and Kenneth D. West, 1994, Automatic lag selection in covariance matrix estimation, *Review of Economic Studies* 61, 631–53.

Pan, Jun, 2002, The jump-risk premia implicit in options: Evidence from an integrated time-series study, *Journal of Financial Economics* 63, 3–50.

Patton, Andrew J, and Kevin Sheppard, 2013, Good volatility, bad volatility: Signed jumps and the persistence of volatility, Working Paper.

Pettenuzzo, Davide, Allan Timmermann, and Rossen Valkanov, 2014, Forecasting stock returns under economic constraints, *Journal of Financial Economics* 114, 517–553.

Prokopczuk, Marcel, and Chardin Wese Simen, 2014, Variance risk premia in commodity markets, Working Paper.

Santa-Clara, Pedro, and Shu Yan, 2010, Crashes, volatility, and the equity premium: Lessons from S&P 500 options, *The Review of Economics and Statistics* 92, 435–451.

Segal, Gill, Ivan Shaliastovich, and Amir Yaron, 2015, Good and bad uncertainty: Macroeconomic and financial market implications, *Journal of Financial Economics* 117, 369–397.

Tsai, Jerry, and Jessica A Wachter, 2015, Rare booms and disasters in a multi-sector endowment economy, *Review of Financial Studies, Forthcoming.*

Weil, Philippe, 1990, Nonexpected utility in macroeconomics, *The Quarterly Journal of Economics* pp. 29–42.

Whitelaw, Robert F, 1994, Time variations and covariations in the expectation and volatility of stock market returns, *The Journal of Finance* 49, 515–541.

Zhou, Hao, 2009, Variance risk premia, asset predictability puzzles, and macroeconomic uncertainty, Working paper.

Tables and Figures

Figure 1: Implied Variance

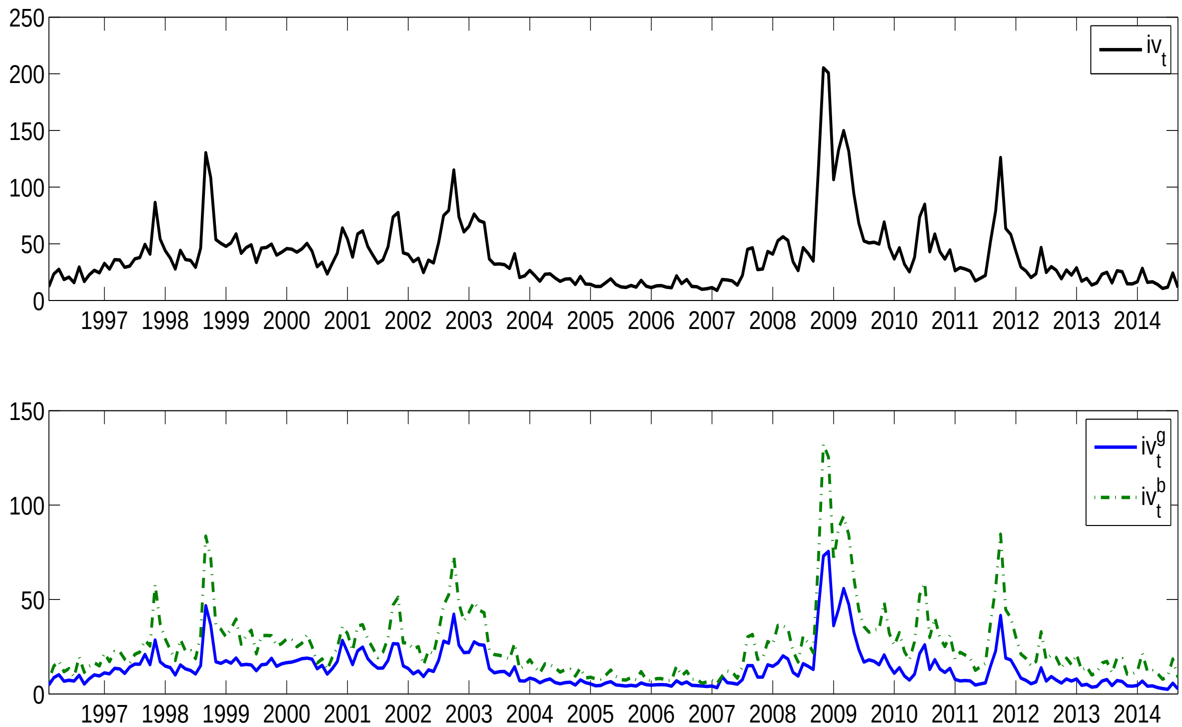

This figure shows the time series of implied variance measures from January 1996 to August 2014. The top panel plots the total implied variance, and the bottom panel depicts good (solid line) and bad (dashed line) implied variances. Measures of implied variance are in units of volatility in percent squared divided by 12.

Figure 2: Realized Variance

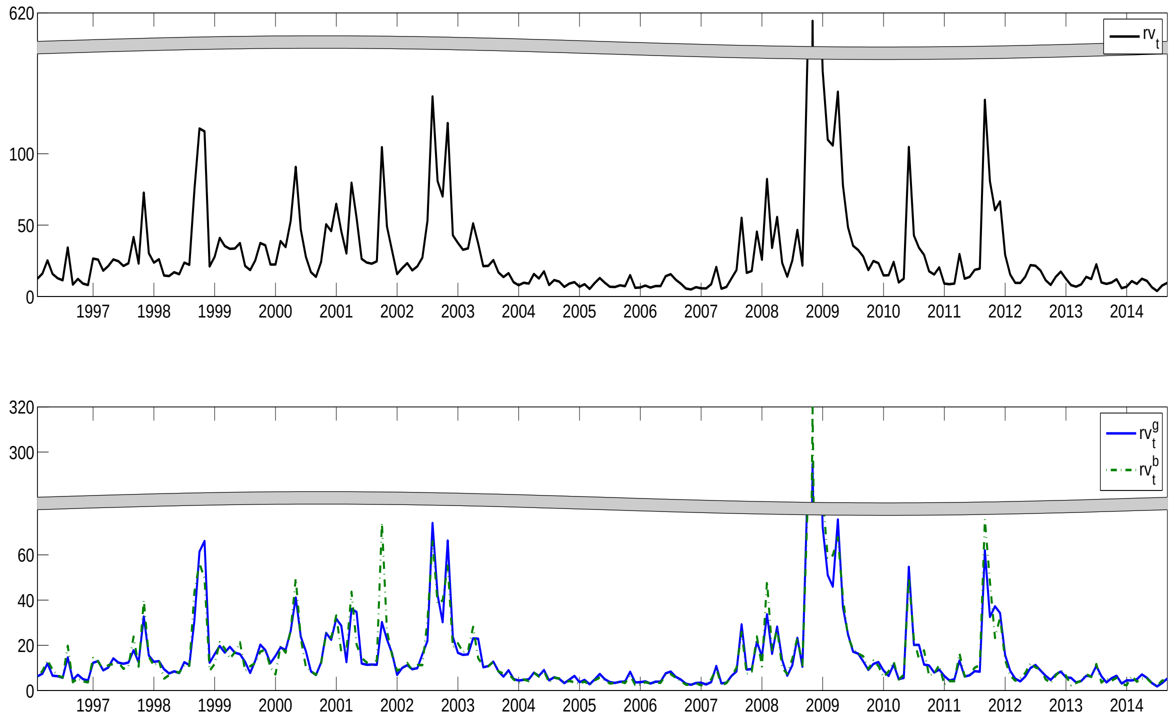

This figure shows the time series of realized variance measures from January 1996 to August 2014. The top panel plots the total realized variance, and the bottom panel depicts good (solid line) and bad (dashed line) realized variances. Measures of realized variance are in units of volatility in percent squared divided by 12.

Figure 3: Variance Premium

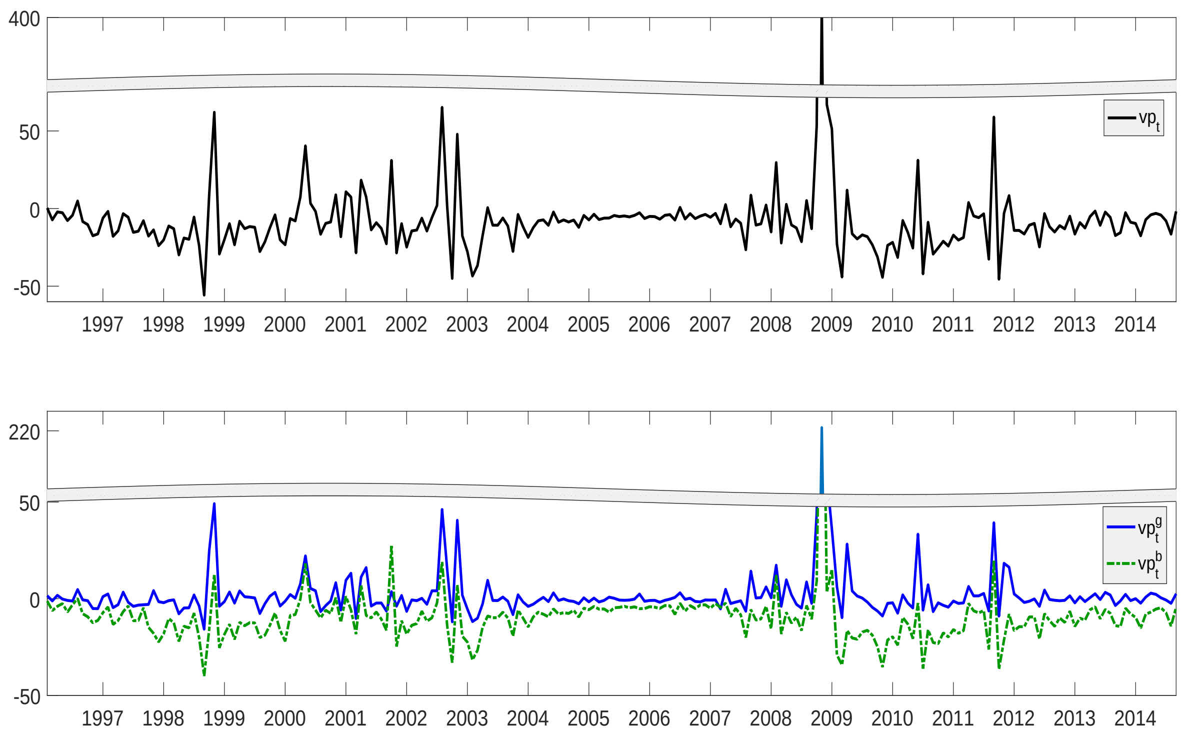

This figure shows the time series of variance premium measures from January 1996 to August 2014. The top panel plots the total variance premium, and the bottom panel depicts good (solid line) and bad (dashed line) variance premia. Measures of variance premium are equal to the difference between the corresponding realized and implied variance measures.

Figure 4: Predicted 1-Year Excess Returns in the Data

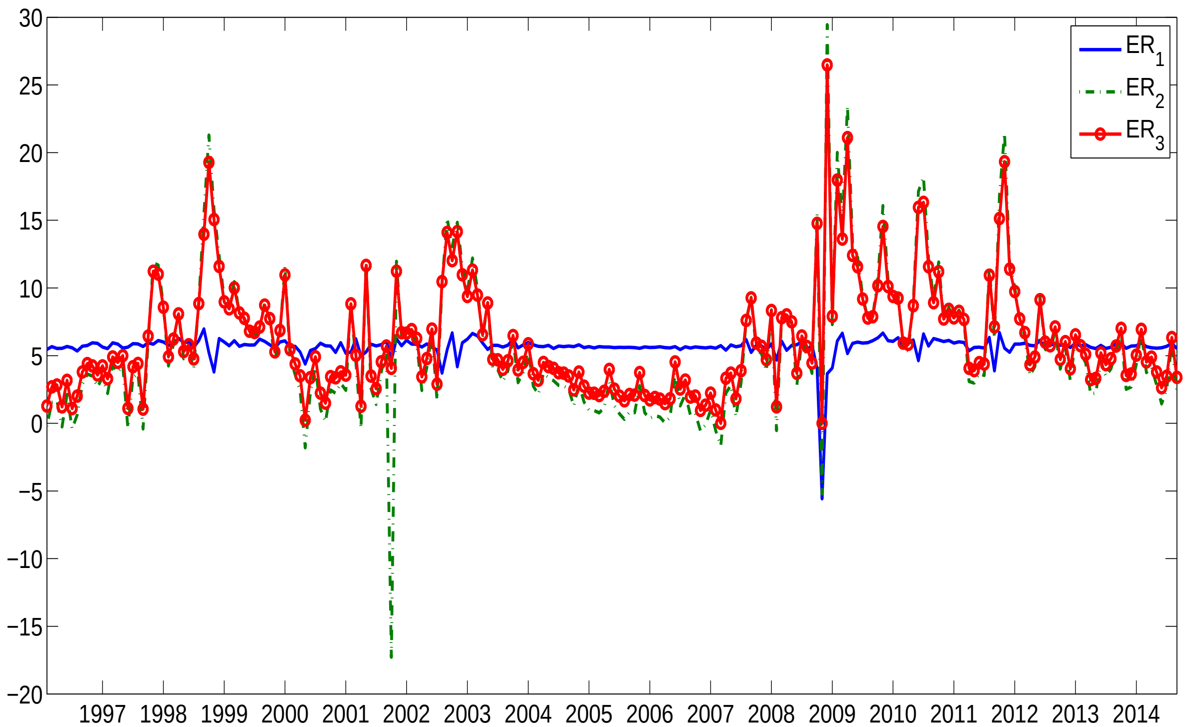

This figure plots the predicted one-year excess aggregate equity returns implied by the predictive regressions. The solid line (ER_1) corresponds to a univariate specification using the total variance premium. The dotted line (ER_2) corresponds to a multivariate specification using the good and bad variance premia jointly. The dashed line (ER_3) corresponds to the multivariate constrained predictability regression using the good and bad variance premia. The values are in annual percentage terms.

Figure 5: Predicted Excess Returns in the Data and Equity Premium in the Model

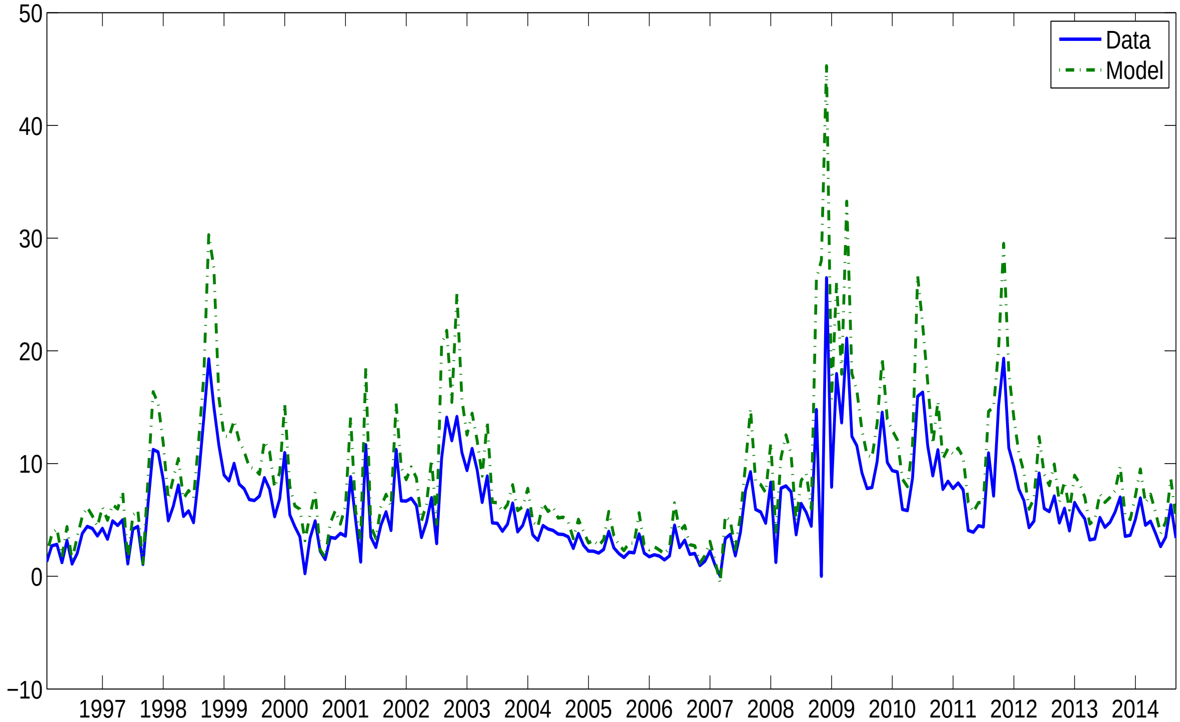

This figure plots the predicted one-year excess aggregate equity returns implied by the constrained regression on the good and bad variance premia (data) and the equity premium implied by the model given the observed good and bad variance premia (model). The values are in annual percentage terms.

Table 1: Summary Statistics for Variance Variables

	Mean	SD	AR(1)
Panel A: Full Sample			
rv_t^g	16.19	25.08	0.62
iv_t^g	13.59	10.74	0.81
vp_t^g	2.60	17.83	0.35
rv_t^b	16.39	26.64	0.57
iv_t^b	25.82	19.06	0.81
vp_t^b	-9.43	16.20	0.16
rv_t	32.58	51.47	0.60
iv_t	39.41	29.60	0.81
vp_t	-6.83	32.87	0.24
Panel B: Excluding the Crisis			
rv_t^g	12.80	12.12	0.57
iv_t^g	12.00	7.63	0.76
vp_t^g	0.79	8.39	0.02
rv_t^b	12.76	12.86	0.54
iv_t^b	23.04	14.04	0.75
vp_t^b	-10.28	8.97	0.14
rv_t	25.56	24.51	0.58
iv_t	35.05	21.37	0.75
vp_t	-9.49	15.58	-0.04

This table shows the summary statistics for realized and implied variance measures, variance premium, and their good and bad components. Panel A presents the results for the benchmark sample from January 1996 to August 2014. Panel B omits the Great Recession period (from December 2007 to June 2009). Measures of implied variance are in units of volatility in percent squared divided by 12.

Table 2: Correlation Matrix of Variance Variables

	iv_t^g	rv_t^g	vp_t^g	iv_t^b	rv_t^b	vp_t^b	iv_t	rv_t
Panel A: Full Sample								
rv_t^g	0.79							
vp_t^g	0.51	0.93						
iv_t^b	0.97	0.78	0.52					
rv_t^b	0.80	0.98	0.90	0.80				
vp_t^b	0.18	0.69	0.87	0.14	0.71			
iv_t	0.99	0.79	0.52	0.99	0.80	0.15		
rv_t	0.80	0.99	0.92	0.80	0.99	0.70	0.80	
vp_t	-0.37	-0.84	0.97	-0.35	-0.83	0.96	-0.36	-0.84
Panel B: Excluding the Crisis								
rv_t^g	0.73							
vp_t^g	0.14	0.78						
iv_t^b	0.95	0.72	0.18					
rv_t^b	0.79	0.93	0.62	0.78				
vp_t^b	-0.35	0.20	0.61	-0.44	0.21			
iv_t	0.98	0.73	0.17	0.99	0.79	-0.41		
rv_t	0.77	0.98	0.71	0.77	0.98	0.21	0.78	
vp_t	0.12	-0.54	0.89	0.16	-0.46	0.90	0.15	-0.51

This table shows the correlations between the realized and implied variance measures, variance premium, and their good and bad components. Panel A presents the results for the benchmark sample from January 1996 to August 2014. Panel B omits the Great Recession period (from December 2007 to June 2009). Measures of implied variance are in units of volatility in percent squared divided by 12.

Table 3: Predicting Aggregate Returns with Variance Premium

	Forecast Horizon h																							
	1				3				6				12				18				24			
	Coef	SE	p-val	$\bar{R}^2$	Coef	SE	p-val	$\bar{R}^2$	Coef	SE	p-val	$\bar{R}^2$	Coef	SE	p-val	$\bar{R}^2$	Coef	SE	p-val	$\bar{R}^2$	Coef	SE	p-val	$\bar{R}^2$
	Panel A: Equity Returns																							
vp_t	-0.35	0.07	0.00	0.04	-0.26	0.05	0.00	0.06	-0.11	0.04	0.00	0.01	-0.03	0.04	0.23	0.00	-0.01	0.04	0.39	0.00	-0.01	0.03	0.38	0.00
		0.23	0.09			0.15	0.05			0.10	0.15			0.06	0.33			0.04	0.41			0.04	0.40	
vp_t^g	-0.04	0.44	0.54	0.04	0.27	0.38	0.24	0.07	0.47	0.23	0.02	0.05	0.52	0.13	0.00	0.07	0.47	0.11	0.00	0.07	0.47	0.13	0.00	0.09
		0.64	0.54			0.50	0.30			0.41	0.12			0.30	0.04			0.26	0.03			0.23	0.02	
vp_t^b	-0.69	0.49	0.08		-0.85	0.45	0.03		-0.75	0.29	0.00		-0.64	0.16	0.00		-0.54	0.14	0.00		-0.54	0.16	0.00	
		0.66	0.15			0.44	0.02			0.36	0.02			0.29	0.01			0.26	0.02			0.24	0.01	
	Panel B: High-Yield Bond Returns																							
vp_t	-0.22	0.09	0.99	0.04	-0.02	0.03	0.74	0.00	0.05	0.03	0.03	0.00	0.07	0.03	0.00	0.04	0.06	0.02	0.00	0.04	0.05	0.02	0.00	0.04
		0.12	0.95			0.08	0.64			0.05	0.16			0.03	0.01			0.02	0.00			0.02	0.00	
vp_t^g	0.36	0.40	0.18	0.07	0.68	0.25	0.00	0.08	0.75	0.23	0.00	0.14	0.61	0.17	0.00	0.20	0.49	0.14	0.00	0.20	0.46	0.13	0.00	0.24
		0.33	0.11			0.26	0.00			0.21	0.00			0.16	0.00			0.13	0.00			0.12	0.00	
vp_t^b	-0.86	0.37	0.01		-0.80	0.30	0.00		-0.72	0.26	0.00		-0.52	0.18	0.00		-0.42	0.15	0.00		-0.40	0.14	0.00	
		0.34	0.01			0.23	0.00			0.19	0.00			0.15	0.00			0.13	0.00			0.12	0.00	
	Panel C: Investment-Grade Bond Returns																							
vp_t	0.09	0.03	0.00	0.02	0.07	0.02	0.00	0.04	0.04	0.01	0.00	0.02	0.05	0.01	0.00	0.06	0.04	0.01	0.00	0.06	0.03	0.01	0.00	0.06
		0.09	0.12			0.05	0.07			0.04	0.11			0.02	0.01			0.02	0.01			0.01	0.00	
vp_t^g	0.45	0.18	0.01	0.05	0.45	0.11	0.00	0.12	0.34	0.11	0.00	0.13	0.28	0.09	0.00	0.18	0.23	0.08	0.00	0.18	0.22	0.07	0.00	0.22
		0.24	0.02			0.19	0.00			0.15	0.01			0.11	0.00			0.10	0.01			0.09	0.00	
vp_t^b	-0.30	0.18	0.05		-0.35	0.13	0.00		-0.29	0.13	0.01		-0.21	0.10	0.02		-0.17	0.09	0.02		-0.17	0.07	0.01	
		0.25	0.09			0.17	0.01			0.13	0.01			0.11	0.02			0.10	0.03			0.09	0.03	

This table presents predictability regression evidence for equity returns (Panel A), high-yield bond returns (Panel B), and investment-grade bond returns (Panel C) from January 1996 to August 2014. The dependent variable is the annualized log excess return over the next h months ($h = 1, 3, 6, 12, 18, 24$). For each asset, the table shows the coefficient (Coef), the standard errors (SE), and the p-values for a univariate predictability regression with the variance premium vp_t, and a multivariate regression with the good and bad variance premia, vp_t^g and vp_t^b, as predictor variables. For each coefficient, the first line reports the Newey and West (1994) standard error with optimal lag selection, and the corresponding p-value. The second line reports the standard error and the p-value from the bootstrap approach. The p-values are calculated for one-sided tests as described in Section 3.4.

Table 4: Predictability in the Cross-Section of Portfolios

		Size						Book-to-Market						Industry					
		3		12		24		3		12		24		3		12		24	
		Coef	$\bar{R}^2$	Coef	$\bar{R}^2$	Coef	$\bar{R}^2$	Coef	$\bar{R}^2$	Coef	$\bar{R}^2$	Coef	$\bar{R}^2$	Coef	$\bar{R}^2$	Coef	$\bar{R}^2$	Coef	$\bar{R}^2$
vp_t	Max	-0.23	0.03	-0.00	0.00	0.01	0.00	-0.18	0.03	0.01	0.00	0.01	0.00	-0.12	0.01	0.01	0.00	0.05	0.00
	Med	-0.28	0.04	-0.02	0.00	-0.00	0.00	-0.28	0.06	-0.02	0.00	-0.00	0.00	-0.21	0.03	-0.02	0.00	-0.01	0.00
	Min	-0.33	0.06	-0.03	0.00	-0.01	0.00	-0.48	0.16	-0.05	0.00	-0.02	0.00	-0.43	0.11	-0.04	0.00	-0.02	0.00
	p-val	0.08		0.52		0.62		0.01		0.42		0.61		0.08		0.71		0.81	
vp_t^g	Max	0.31	0.04	0.81	0.05	0.72	0.07	0.71	0.04	0.74	0.01	0.59	0.06	1.05	0.01	1.14	0.00	0.84	0.02
	Med	0.24	0.05	0.63	0.08	0.63	0.19	-0.02	0.07	0.37	0.03	0.43	0.09	0.16	0.06	0.40	0.05	0.44	0.08
	Min	0.19	0.08	0.48	0.10	0.42	0.23	-0.54	0.15	0.20	0.10	0.34	0.15	-0.43	0.11	-0.02	0.11	0.22	0.15
	p-val	0.47		0.04		0.02		0.20		0.04		0.03		0.19		0.02		0.03	
vp_t^b	Max	-0.77		-0.57		-0.49		-0.34		-0.30		-0.40		0.07		-0.03		-0.29	
	Med	-0.83		-0.74		-0.69		-0.62		-0.45		-0.50		-0.63		-0.48		-0.48	
	Min	-0.94		-0.91		-0.80		-1.25		-0.86		-0.65		-1.63		-1.30		-0.82	
	p-val	0.09		0.02		0.02		0.01		0.01		0.02		0.02		0.01		0.04	

This table presents predictability regression evidence for a cross-section of equity returns from January 1996 to August 2014. The dependent variable is the annualized log excess return over the next h months ($h = 1, 3, 6, 12, 18, 24$). The table shows the minimum, median, and maximum values of the coefficient (Coef) and the joint p-values for a univariate predictability regression with the variance premium vp_t, and a multivariate regression with the good and bad variance premia, vp_t^g and vp_t^b, as predictor variables among ten size, book-to-market, and industry portfolios. The joint p-values are computed as described in Appendix A.2.

Table 5: Predicting Aggregate Returns with Variance Premium: Alternative Expected Variance Measure

	Forecast Horizon h																							
	1				3				6				12				18				24			
	Coef	SE	p-val	$\bar{R}^2$	Coef	SE	p-val	$\bar{R}^2$	Coef	SE	p-val	$\bar{R}^2$	Coef	SE	p-val	$\bar{R}^2$	Coef	SE	p-val	$\bar{R}^2$	Coef	SE	p-val	$\bar{R}^2$
Panel A: Equity Returns																								
vp_t	-0.18	0.44	0.34	0.00	-0.29	0.34	0.20	0.01	-0.39	0.18	0.01	0.05	-0.26	0.11	0.01	0.04	-0.26	0.08	0.00	0.06	-0.30	0.09	0.00	0.11
		0.40	0.33			0.32	0.17			0.26	0.06			0.19	0.08			0.16	0.05			0.14	0.01	
vp_t^g	3.61	0.89	0.00	0.05	3.72	0.54	0.00	0.13	2.55	0.43	0.00	0.13	1.39	0.39	0.00	0.08	1.27	0.37	0.00	0.11	1.35	0.35	0.00	0.17
		1.71	0.02			1.20	0.00			0.83	0.00			0.63	0.01			0.57	0.01			0.51	0.00	
vp_t^b	-1.13	0.47	0.01		-1.43	0.32	0.00		-1.21	0.24	0.00		-0.76	0.20	0.00		-0.72	0.16	0.00		-0.79	0.18	0.00	
		0.68	0.04			0.54	0.00			0.42	0.00			0.35	0.01			0.32	0.01			0.30	0.00	
Panel B: High-Yield Bond Returns																								
vp_t	-0.28	0.32	0.19	0.01	-0.53	0.17	0.00	0.13	-0.54	0.13	0.00	0.22	-0.38	0.11	0.00	0.23	-0.34	0.07	0.00	0.27	-0.31	0.06	0.00	0.31
		0.21	0.06			0.16	0.00			0.13	0.00			0.10	0.00			0.08	0.00			0.07	0.00	
vp_t^g	2.65	0.96	0.00	0.07	1.80	0.51	0.00	0.14	1.14	0.42	0.00	0.20	0.49	0.29	0.05	0.24	0.52	0.23	0.01	0.28	0.54	0.22	0.01	0.34
		0.91	0.01			0.63	0.01			0.44	0.01			0.33	0.07			0.30	0.04			0.27	0.02	
vp_t^b	-0.96	0.39	0.01		-1.12	0.31	0.00		-0.95	0.27	0.00		-0.61	0.20	0.00		-0.57	0.15	0.00		-0.55	0.14	0.00	
		0.35	0.00			0.28	0.00			0.22	0.00			0.18	0.00			0.17	0.00			0.15	0.00	
Panel C: Investment-Grade Bond Returns																								
vp_t	-0.28	0.08	0.00	0.04	-0.26	0.07	0.00	0.11	-0.22	0.07	0.00	0.15	-0.19	0.05	0.00	0.22	-0.17	0.04	0.00	0.27	-0.15	0.03	0.00	0.28
		0.15	0.02			0.12	0.01			0.10	0.01			0.07	0.00			0.06	0.00			0.05	0.00	
vp_t^g	-0.03	0.36	0.54	0.05	0.08	0.24	0.37	0.13	0.21	0.22	0.18	0.16	0.07	0.15	0.31	0.24	0.11	0.13	0.20	0.27	0.11	0.11	0.16	0.30
		0.66	0.49			0.46	0.41			0.32	0.24			0.24	0.36			0.21	0.30			0.19	0.28	
vp_t^b	-0.32	0.16	0.02		-0.35	0.14	0.01		-0.33	0.14	0.01		-0.25	0.10	0.00		-0.23	0.08	0.00		-0.21	0.07	0.00	
		0.26	0.08			0.21	0.03			0.16	0.01			0.13	0.02			0.12	0.02			0.11	0.02	

This table presents predictability regression evidence for equity returns (Panel A), high-yield bond returns (Panel B), and investment-grade bond returns (Panel C) from January 1996 to August 2014. The dependent variable is the annualized log excess return over the next h months ($h = 1, 3, 6, 12, 18, 24$). The variance premia measurements are based on the ex-ante physical variances constructed from the fitted expectations of the next month squared returns. For each asset, the table shows the coefficient (Coef), the standard errors (SE), and the p-values for a univariate predictability regression with the variance premium vp_t, and a multivariate regression with the good and bad variance premia, vp_t^g and vp_t^b, as predictor variables. For each coefficient, the first line reports the Newey and West (1994) standard error with optimal lag selection, and the corresponding p-value. The second line reports the standard error and the p-value from the bootstrap approach. The p-values are calculated for one-sided tests as described in Section 3.4.

Table 6: Predicting Aggregate Returns with Variance Premium: Alternative standard errors

	Forecast Horizon h																	
	1			3			6			12			18			24		
	Coef	p-val	$\bar{R}^2$	Coef	p-val	$\bar{R}^2$	Coef	p-val	$\bar{R}^2$	Coef	p-val	$\bar{R}^2$	Coef	p-val	$\bar{R}^2$	Coef	p-val	$\bar{R}^2$
Panel A: Equity Returns																		
vp_t	-0.35	0.04	0.04	-0.26	0.01	0.06	-0.11	0.14	0.01	-0.03	0.32	0.00	-0.01	0.41	0.00	-0.01	0.40	0.00
		0.00			0.00			0.03			0.30			0.45			0.48	
vp_t^g	-0.04	0.53	0.04	0.27	0.27	0.07	0.47	0.11	0.05	0.52	0.02	0.07	0.47	0.02	0.07	0.47	0.01	0.09
		0.45			0.42			0.07			0.00			0.00			0.00	
vp_t^b	-0.69	0.12		-0.85	0.02		-0.75	0.01		-0.64	0.01		-0.54	0.01		-0.54	0.01	
		0.07			0.11			0.04			0.00			0.00			0.00	
Panel B: High-Yield Bond Returns																		
vp_t	-0.22	0.83	0.04	-0.02	0.57	0.00	0.05	0.26	0.00	0.07	0.06	0.04	0.02	0.03	0.04	0.05	0.02	0.04
		0.96			0.74			0.00			0.00			0.00			0.00	
vp_t^g	0.36	0.25	0.07	0.68	0.02	0.08	0.75	0.00	0.14	0.61	0.00	0.20	0.49	0.00	0.20	0.46	0.00	0.24
		0.11			0.00			0.00			0.00			0.00			0.00	
vp_t^b	-0.86	0.04		-0.80	0.00		-0.72	0.00		-0.52	0.00		-0.42	0.00		-0.40	0.00	
		0.00			0.00			0.00			0.00			0.00			0.00	
Panel C: Investment-Grade Bond Returns																		
vp_t	0.09	0.13	0.02	0.07	0.13	0.04	0.04	0.12	0.02	0.05	0.02	0.06	0.04	0.01	0.06	0.03	0.01	0.06
		0.00			0.02			0.00			0.00			0.00			0.00	
vp_t^g	0.45	0.06	0.05	0.45	0.01	0.12	0.34	0.01	0.13	0.28	0.00	0.18	0.23	0.00	0.18	0.22	0.00	0.22
		0.00			0.00			0.00			0.01			0.01			0.00	
vp_t^b	-0.30	0.13		-0.35	0.01		-0.29	0.01		-0.21	0.01		-0.17	0.01		-0.17	0.01	
		0.01			0.01			0.01			0.04			0.05			0.03	

This table presents predictability regression evidence for equity returns (Panel A), high-yield bond returns (Panel B), and investment-grade bond returns (Panel C) from January 1996 to August 2014. The dependent variable is the annualized log excess return over the next h months ($h = 1, 3, 6, 12, 18, 24$). For each asset, the table shows the coefficient (Coef) and the one-sided p-values for a univariate predictability regression with the variance premium vp_t, and a multivariate regression with the good and bad variance premia, vp_t^g and vp_t^b, as predictor variables. For each coefficient, the first line reports the inference based on Hodrick (1992) standard errors, while the second line reports the inference based on the ARIMA(1,1)-GARCH(1,1) parametric bootstrap approach of Bakshi, Panayotov, and Skoulakis (2011).

Table 7: Predicting Excess Returns with Variance Premium in a Longer Sample

	Forecast Horizon h																							
	1				3				6				12				18				24			
	Coef	SE	p-val	$\bar{R}^2$	Coef	SE	p-val	$\bar{R}^2$	Coef	SE	p-val	$\bar{R}^2$	Coef	SE	p-val	$\bar{R}^2$	Coef	SE	p-val	$\bar{R}^2$	Coef	SE	p-val	$\bar{R}^2$
Panel A: Equity Returns																								
vp_t	-0.33	0.07	0.00	0.03	-0.26	0.05	0.00	0.05	-0.10	0.03	0.00	0.01	-0.03	0.03	0.20	0.00	-0.01	0.03	0.37	0.00	-0.01	0.02	0.39	0.00
		0.23	0.09			0.15	0.05			0.10	0.16			0.06	0.31			0.04	0.40			0.04	0.43	
vp_t^g	-0.13	0.38	0.63	0.03	0.14	0.31	0.32	0.06	0.30	0.21	0.08	0.03	0.36	0.12	0.00	0.04	0.32	0.10	0.00	0.04	0.27	0.11	0.01	0.04
		0.55	0.61			0.44	0.38			0.36	0.21			0.28	0.10			0.24	0.09			0.22	0.11	
vp_t^b	-0.56	0.40	0.08		-0.70	0.36	0.03		-0.55	0.25	0.01		-0.46	0.14	0.00		-0.37	0.12	0.00		-0.31	0.13	0.01	
		0.56	0.16			0.38	0.03			0.31	0.04			0.26	0.04			0.24	0.06			0.22	0.08	
Panel B: High-Yield Bond Returns																								
vp_t	-0.19	0.09	0.98	0.03	-0.02	0.03	0.71	0.00	0.05	0.03	0.02	0.00	0.07	0.03	0.00	0.03	0.06	0.02	0.00	0.03	0.06	0.01	0.00	0.03
		0.12	0.94			0.08	0.64			0.05	0.16			0.03	0.01			0.02	0.00			0.02	0.00	
vp_t^g	0.16	0.36	0.33	0.03	0.52	0.23	0.01	0.05	0.63	0.21	0.00	0.10	0.52	0.16	0.00	0.14	0.42	0.13	0.00	0.14	0.35	0.12	0.00	0.14
		0.29	0.30			0.23	0.01			0.19	0.00			0.15	0.00			0.13	0.00			0.12	0.00	
vp_t^b	-0.59	0.34	0.04		-0.62	0.27	0.01		-0.58	0.24	0.01		-0.42	0.18	0.01		-0.33	0.14	0.01		-0.27	0.13	0.02	
		0.30	0.02			0.20	0.00			0.17	0.00			0.14	0.00			0.13	0.00			0.12	0.01	
Panel C: Investment-Grade Bond Returns																								
vp_t	0.10	0.02	0.00	0.02	0.07	0.02	0.00	0.03	0.04	0.01	0.00	0.02	0.05	0.02	0.00	0.04	0.04	0.01	0.00	0.05	0.03	0.01	0.00	0.05
		0.09	0.12			0.06	0.09			0.04	0.14			0.02	0.02			0.02	0.02			0.01	0.01	
vp_t^g	0.29	0.16	0.03	0.03	0.31	0.11	0.00	0.06	0.25	0.11	0.02	0.07	0.22	0.09	0.01	0.11	0.19	0.07	0.01	0.13	0.17	0.06	0.00	0.15
		0.22	0.09			0.18	0.03			0.15	0.04			0.11	0.02			0.10	0.02			0.09	0.02	
vp_t^b	-0.11	0.16	0.25		-0.19	0.13	0.07		-0.18	0.13	0.08		-0.15	0.10	0.06		-0.13	0.08	0.05		-0.12	0.06	0.03	
		0.23	0.30			0.16	0.10			0.13	0.06			0.11	0.07			0.10	0.09			0.09	0.09	

This table presents predictability regression evidence for equity returns (Panel A), high-yield bond returns (Panel B), and investment-grade bond returns (Panel C) from January 1988 to August 2014. The dependent variable is the annualized log excess return over the next h months ($h = 1, 3, 6, 12, 18, 24$). For each asset, the table shows the coefficient (Coef), the standard errors (SE), and the p-values for a univariate predictability regression with the variance premium vp_t, and a multivariate regression with the good and bad variance premia, vp_t^g and vp_t^b, as predictor variables. For each coefficient, the first line reports the Newey and West (1994) standard error with optimal lag selection, and the corresponding p-value. The second line reports the standard error and the p-value from the bootstrap approach. The p-values are calculated for one-sided tests as described in Section 3.4.

Table 8: Predicting Excess Returns with Variance Premium Excluding the Great Recession

	Forecast Horizon h																							
	1				3				6				12				18				24			
	Coef	SE	p-val	$\bar{R}^2$	Coef	SE	p-val	$\bar{R}^2$	Coef	SE	p-val	$\bar{R}^2$	Coef	SE	p-val	$\bar{R}^2$	Coef	SE	p-val	$\bar{R}^2$	Coef	SE	p-val	$\bar{R}^2$
Panel A: Equity Returns																								
vp_t	-0.39	0.29	0.09	0.01	-0.28	0.14	0.02	0.02	-0.20	0.10	0.02	0.02	-0.14	0.07	0.03	0.01	-0.13	0.08	0.06	0.01	-0.12	0.07	0.04	0.01
		0.31	0.11			0.16	0.04			0.11	0.04			0.08	0.04			0.06	0.01			0.06	0.02	
vp_t^g	0.01	0.65	0.49	0.01	0.50	0.43	0.13	0.05	0.33	0.27	0.11	0.05	0.37	0.14	0.00	0.06	0.34	0.17	0.02	0.06	0.33	0.19	0.04	0.07
		0.70	0.49			0.51	0.15			0.39	0.18			0.28	0.08			0.24	0.07			0.21	0.05	
vp_t^b	-0.76	0.45	0.05		-1.00	0.36	0.00		-0.69	0.27	0.01		-0.61	0.19	0.00		-0.57	0.19	0.00		-0.53	0.23	0.01	
		0.64	0.11			0.47	0.01			0.39	0.03			0.33	0.03			0.29	0.02			0.26	0.02	
Panel B: High-Yield Bond Returns																								
vp_t	0.03	0.16	0.42	0.00	-0.07	0.07	0.83	0.00	-0.02	0.07	0.61	0.00	-0.01	0.05	0.54	0.00	-0.02	0.05	0.67	0.00	-0.02	0.04	0.67	0.00
		0.15	0.42			0.08	0.83			0.05	0.64			0.04	0.56			0.03	0.78			0.03	0.73	
vp_t^g	0.44	0.28	0.06	0.01	0.44	0.14	0.00	0.06	0.40	0.17	0.01	0.07	0.31	0.18	0.04	0.06	0.24	0.17	0.08	0.04	0.26	0.16	0.05	0.07
		0.33	0.08			0.24	0.02			0.19	0.01			0.14	0.01			0.12	0.02			0.10	0.00	
vp_t^b	-0.34	0.25	0.09		-0.54	0.21	0.01		-0.40	0.19	0.02		-0.30	0.18	0.05		-0.26	0.17	0.06		-0.28	0.18	0.07	
		0.31	0.12			0.23	0.01			0.18	0.01			0.16	0.03			0.14	0.03			0.13	0.01	
Panel C: Investment-Grade Bond Returns																								
vp_t	0.06	0.07	0.22	0.00	0.02	0.03	0.32	0.00	0.03	0.03	0.18	0.00	0.00	0.02	0.49	-0.01	0.00	0.02	0.45	-0.01	-0.01	0.02	0.63	0.00
		0.11	0.33			0.06	0.40			0.04	0.27			0.03	0.49			0.02	0.44			0.02	0.63	
vp_t^g	0.32	0.20	0.05	0.01	0.31	0.11	0.00	0.05	0.20	0.13	0.06	0.04	0.11	0.12	0.18	0.02	0.09	0.10	0.18	0.02	0.12	0.09	0.09	0.05
		0.25	0.09			0.18	0.03			0.14	0.06			0.10	0.13			0.09	0.13			0.08	0.06	
vp_t^b	-0.18	0.19	0.16		-0.25	0.12	0.02		-0.14	0.11	0.11		-0.10	0.12	0.21		-0.08	0.11	0.23		-0.12	0.10	0.12	
		0.23	0.19			0.17	0.05			0.14	0.15			0.12	0.19			0.11	0.22			0.10	0.10	

This table presents predictability regression evidence for equity returns (Panel A), high-yield bond returns (Panel B), and investment-grade bond returns (Panel C) from January 1996 to August 2014, omitting the Great Recession (from December 2007 to June 2009). The dependent variable is the annualized log excess return over the next h months ($h = 1, 3, 6, 12, 18, 24$). For each asset, the table shows the coefficient (Coef), the standard errors (SE), and the p-values for a univariate predictability regression with the variance premium vp_t, and a multivariate regression with the good and bad variance premia, vp_t^g and vp_t^b, as predictor variables. For each coefficient, the first line reports the Newey and West (1994) standard error with optimal lag selection, and the corresponding p-value. The second line reports the standard error and the p-value from the bootstrap approach. The p-values are calculated for one-sided tests as described in Section 3.4.

Table 9: Predicting Excess Returns with Variance Premium with Constrained Regression

	Forecast Horizon h																							
	1				3				6				12				18				24			
	Coef	SE	p-val	$\bar{R}^2$	Coef	SE	p-val	$\bar{R}^2$	Coef	SE	p-val	$\bar{R}^2$	Coef	SE	p-val	$\bar{R}^2$	Coef	SE	p-val	$\bar{R}^2$	Coef	SE	p-val	$\bar{R}^2$
	Panel A: Equity Returns																							
vp_t	-0.02	0.24	0.47	0.00	-0.02	0.17	0.46	0.00	-0.01	0.05	0.39	0.00	-0.01	0.03	0.30	0.00	-0.01	0.03	0.40	0.00	-0.01	0.03	0.37	0.00
		0.08	0.18			0.07	0.17			0.05	0.18			0.04	0.17			0.03	0.28			0.02	0.26	
vp_t^g	0.37	0.67	0.29	0.00	0.53	0.54	0.17	0.03	0.50	0.24	0.02	0.05	0.44	0.18	0.01	0.06	0.41	0.16	0.01	0.06	0.39	0.19	0.02	0.09
		0.34	0.33			0.30	0.14			0.26	0.11			0.21	0.07			0.19	0.06			0.17	0.04	
vp_t^b	-0.45	0.46	0.16		-0.63	0.38	0.05		-0.59	0.28	0.02		-0.52	0.21	0.01		-0.48	0.19	0.00		-0.46	0.22	0.02	
		0.33	0.17			0.28	0.05			0.24	0.03			0.21	0.03			0.20	0.03			0.19	0.02	
	Panel B: High-Yield Bond Returns																							
vp_t	-0.01	0.22	0.52	0.00	-0.01	0.03	0.66	0.00	0.05	0.03	0.04	0.00	0.08	0.02	0.00	0.04	0.06	0.02	0.00	0.04	0.05	0.02	0.00	0.04
		0.05	0.87			0.04	0.86			0.03	0.16			0.02	0.00			0.02	0.00			0.01	0.00	
vp_t^g	0.46	0.49	0.17	0.03	0.48	0.28	0.04	0.07	0.52	0.30	0.04	0.13	0.48	0.25	0.03	0.19	0.44	0.18	0.01	0.21	0.42	0.17	0.01	0.25
		0.20	0.05			0.16	0.02			0.14	0.00			0.11	0.00			0.10	0.00			0.09	0.00	
vp_t^b	-0.55	0.41	0.09		-0.56	0.33	0.04		-0.48	0.33	0.07		-0.38	0.27	0.08		-0.36	0.20	0.04		-0.36	0.19	0.03	
		0.20	0.01			0.16	0.00			0.14	0.00			0.12	0.00			0.11	0.00			0.10	0.00	
	Panel C: Investment-Grade Bond Returns																							
vp_t	0.08	0.03	0.01	0.02	0.07	0.02	0.00	0.04	0.04	0.01	0.00	0.03	0.05	0.01	0.00	0.06	0.04	0.01	0.00	0.06	0.03	0.01	0.00	0.06
		0.03	0.06			0.03	0.03			0.02	0.09			0.01	0.01			0.01	0.01			0.01	0.00	
vp_t^g	0.39	0.17	0.01	0.05	0.37	0.12	0.00	0.12	0.33	0.14	0.01	0.13	0.30	0.13	0.01	0.19	0.25	0.10	0.01	0.20	0.24	0.09	0.00	0.24
		0.13	0.02			0.11	0.01			0.10	0.01			0.08	0.00			0.07	0.00			0.06	0.00	
vp_t^b	-0.24	0.17	0.08		-0.27	0.15	0.03		-0.28	0.16	0.04		-0.23	0.14	0.04		-0.20	0.12	0.05		-0.20	0.09	0.02	
		0.13	0.08			0.11	0.03			0.09	0.01			0.08	0.01			0.08	0.02			0.07	0.01	

This table presents predictability regression evidence for equity returns (Panel A), high-yield bond returns (Panel B), and investment-grade bond returns (Panel C) from January 1996 to August 2014 using the constrained regression approach described in Section 3.5. The dependent variable is the annualized log excess return over the next h months ($h = 1, 3, 6, 12, 18, 24$). For each asset, the table shows the coefficient (Coef), the standard errors (SE), and the p-values for a univariate predictability regression with the variance premium vp_t, and a multivariate regression with the good and bad variance premia, vp_t^g and vp_t^b, as predictor variables. For each coefficient, the first line reports the Newey and West (1994) standard error with optimal lag selection, and the corresponding p-value. The second line reports the standard error and the p-value from the bootstrap approach. The p-values are calculated for one-sided tests as described in Section 3.4.