
Short-Run Bond Risk Premia

Philippe Mueller

*Finance Group, Warwick Business School
Gibbet Hill Road, Coventry CV4 7AL, UK
philippe.mueller@wbs.ac.uk*

Andrea Vedolin

*Questrom School of Business, Boston University
595 Commonwealth Avenue, Boston 02215, MA, USA
avedolin@bu.edu*

Hao Zhou*

*PBC School of Finance, Tsinghua University
43 Chengfu Road, Haidian District
Beijing 100083, P. R. China
zhouh@pbcfs.tsinghua.edu.cn*

Published 12 June 2019

In the short-run, bond risk premia exhibit pronounced spikes around major economic and financial crises. In contrast, long-term bond risk premia feature cyclical swings. We empirically examine the predictability of the market variance risk premium — a proxy of economic uncertainty — for bond risk premia and we show the strong predictive power for the one-month horizon that quickly recedes for longer horizons. The variance risk premium is largely orthogonal to well-established bond return predictors — forward rates, jumps, and macro variables. We rationalize our empirical findings in an equilibrium model of uncertainty about consumption and inflation which is coupled with recursive preferences. We show that the model can quantitatively explain the levels of bond and variance risk premia as well as the predictive power of the variance risk premium, while jointly matching salient features of other asset prices.

Keywords: Variance risk premium; bond risk premia; expectations hypothesis; inflation dynamics; economic uncertainty.

JEL Classifications: E43, E44, G12, G13, G17

*Corresponding author.

1. Introduction

The failure of the expectations hypothesis of the term structure of interest rates, first documented in Fama and Bliss (1987) and Campbell and Shiller (1991), has received unprecedented attention in both empirical and theoretical academic literature over the past 20 years. In this paper, we first document the large and significant predictive power of the variance risk premium, defined as the difference between the risk-neutral and statistical expectations of realized variance, for bond risk premia at very short horizons. This short-run forecastability is orthogonal to the well-documented long horizon predictability from forward rates (Cochrane and Piazzesi, 2005), macro variables (Ludvigson and Ng, 2009), and jump risk (Wright and Zhou, 2009).¹ We then posit an economy with time-varying uncertainty risk about real and nominal quantities coupled with agents' preferences for an early resolution of uncertainty, and show that these ingredients are enough to quantitatively explain the violation of the expectation hypothesis while matching the moments of the variance risk premium, equity premium, and risk-free rate.

To capture this short-run uncertainty component of bond risk premia, we rely on the market variance risk premium — or the difference between risk-neutral and objective expectations of the return variation. Following the path of previous work, we proxy the risk-neutral expected variance by the popular VIX² index, which is termed as the “market gauge of fear” (Whaley, 2000). With high-frequency intraday data of futures on the S&P 500, we use heterogeneous autoregressive models of realized variance (HAR-RV model, see Corsi, 2009) augmented by lags of implied variances (Drechsler and Yaron, 2011) for estimating the objective expectation of variance risk. Our average variance risk premium is 21.57 (percentage squared monthly basis) and it falls within the typical range of recent empirical estimates. More importantly, our time-series of variance risk premium always remains positive, which makes it a natural candidate for measuring economic uncertainty or even stochastic risk aversion.

We document the predictive power of the bond variance risk premium for short-run bond risk premia using various data. We show that the variance

¹Cieslak and Povala (2010) decompose long-term yields into a persistent component and cycles and find that the cyclical component is a strong predictor encompassing several other ones. Duffee (2011) estimates a five factor Gaussian model using Treasury yields and extracts a latent factor, that is “hidden” from — or weakly spanned by — the cross-section of yields but has bearing on excess bond returns. Huang and Shi (2010) construct a single macro factor using a group lasso method and show that this factor almost doubles the R^2 compared to Ludvigson and Ng (2009).

risk premium is a significant predictor for one-month excess returns on bond portfolios, with underlying maturities ranging between zero and 10 years, obtained from CRSP. The same results hold when calculating one-month excess returns on bonds with maturities ranging from two months to 10 years, calculated using the [Gürkaynak *et al.* \(2007\)](#) dataset. The same results hold when using the variance risk premium to forecast one-month Treasury bill excess returns. However, the variance risk premium only has negligible forecasting power for longer horizon excess returns and in particular, it has zero predictive power for one-year excess returns on two to five year Treasury bonds, which are in general used to run bond predictability regressions. We show that the short-run forecasting power is robust to the inclusion of other well-established bond risk premium predictors, such as forward rates, macro variables, and jump risk. While these variables have previously been shown to predict bond risk premia for longer maturities, they do not subsume the significance of the variance risk premium at shorter horizons and in some cases even have zero predictive power.

The intuition for our empirical result becomes more evident when we look at the time series of short-term bond risk premia. Bond risk premia at short horizons exhibit pronounced spikes around major economic and financial crises. This pattern is distinctly different from the cyclical swings with a length of up to several years, typically observed in long-term bond risk premia (see [Fama and Bliss, 1987](#); [Cochrane and Piazzesi, 2005](#)). Interestingly, the variance risk premium exhibits a similar time-series behavior as the short-term bonds: It rises sharply before economic or financial crises and then drops again. On the other hand, standard predictors like the CP factor display a strong cyclical behavior (see [Kojien *et al.*, 2010](#)). The upshot is that short-term variation in bond risk premia is related to economic uncertainty, which is short-lived (see [Bloom, 2009](#)), as opposed to a business cycle component, which is more apparent in bond risk premia of longer maturities.

We propose a potential explanation for this short-run predictability in an economy with time-varying economic uncertainty about real and nominal quantities, extending the real uncertainty model of [Bollerslev *et al.* \(2009\)](#).² In an economy with stochastic inflation volatility but with only exogenous shocks, money neutrality holds and there is no inflation risk premium except for the standard Jensen's inequality term (see [Zhou, 2010](#)). In this model with

² [Wu \(2008\)](#), [Hasseltoft \(2010\)](#), and [Doh \(2010\)](#) study the long-run risk models for term structure with both real and nominal uncertainty. However, they also rely on the small persistent growth component and they do not examine the predictability of the variance risk premium.

endogenous inflation shocks, we derive a genuine inflation risk premium through two channels. First, we introduce an endogenous stochastic volatility process through the consumption growth channel. Second, we let the stochastic volatility process be correlated with the consumption uncertainty channel. While the equilibrium model developed in this paper is related to the long-run risk model of [Bansal and Shaliastovich \(2010\)](#), we explicitly abstain from modeling the small persistent component in consumption growth and inflation as done in their setup. In our model, we allow the volatility of both inflation and consumption — or the economic uncertainty about these quantities — to speak for themselves on how far the model can go to accommodate the observed level and predictability in bond risk premia.

The key to matching the bond risk premium dynamics is through the calibrating the inflation process, while leaving the choices of preference parameters and real economy dynamics similar to existing studies (see, e.g., [Bansal and Yaron, 2004](#), [Bollerslev et al., 2009](#)). Our calibration exercise shows that an autonomous inflation process (with or without stochastic volatility) is not able to replicate the size of the bond risk premium. Combining both a consumption growth channel and an uncertainty channel of non-neutral inflation dynamics, leads to reasonable and rich bond risk premia. Indeed, our calibrated numbers are only several basis points away from their empirical counterparts. We also show that the model produces a reasonable equity premium and a risk-free rate but overshoots the risk-free rate volatility. While the higher order moments (kurtosis and skewness) of the variance risk premium are fitted quite well, the average variance risk premium produced by our model is slightly smaller than its empirical estimate. Finally, the predictive power of the equity variance risk premium for bond risk premia is fitted remarkably well by our preferred inflation uncertainty model.

Previous work has attempted to explain the failure of the expectations hypothesis through the growth channel of consumption, e.g., [Wachter \(2006\)](#) (external habit), [Bansal and Shaliastovich \(2010\)](#) (long-run risk), [Gabaix \(2009\)](#) (rare disasters), [Xiong and Yan \(2010\)](#) (heterogeneous expectations), and [Vayanos and Vila \(2009\)](#) (preferred habitat). We argue that adding the inflation uncertainty component can go a long way in fitting the salient features of asset prices and bond risk premia in particular.³ We also

³Papers that study the impact of frictions in bond markets on bond returns include [Greenwood and Vayanos \(2010\)](#) (bond supply) and [Fontaine and Garcia \(2010\)](#) (liquidity premium). [Buraschi and Whelan \(2010\)](#) study the impact of dispersion in forecasts on economic quantities on bond returns and estimate highly significant coefficients.

contribute to the growing macroeconomic literature that emphasizes the quantitative importance of time-varying volatility in real and nominal variables to understand the source of aggregate fluctuations, the evolution of the economy, and policy analysis (see, e.g., [Fernández-Villaverde and Rubio-Ramírez, 2010](#)). Similarly, [Bloom \(2009\)](#) and [Bloom *et al.* \(2010\)](#) show that higher economic uncertainty, proxied by the VIX, decreases employment and output in near terms. Our empirical finding and modeling approach are broadly consistent with the macroeconomic uncertainty framework driven by real and nominal volatility dynamics.

The rest of the paper is organized as follows. Section 2 describes our data set and the methods used to estimate the variance risk premium and provides empirical finding for the bond return predictability of the variance risk premium. Section 3 presents a structural model of inflation uncertainty with calibration evidence for risk premium dynamics. Section 4 concludes.

2. Empirical Analysis

In this section, we first discuss the data we use in our empirical analysis — excess returns on Treasury bond portfolios, T-bills and Treasury notes, macroeconomic and financial variables, daily VIX levels, and high-frequency S&P 500 index returns. We measure the equity market variance risk premium as the difference between the squared VIX values and a forecast of realized variance, using a heterogeneous autoregressive forecasting equation augmented by multiple lags of implied variances.⁴ We then present evidence for the predictive power of the variance risk premium for bond risk premia at short horizons. We first run a set of univariate regressions using the variance risk premium as the sole predictor variable and then control for other well-established predictors. We find that the equity variance risk premium is a robust predictor for bond risk premia at short horizons but has very limited predictive power for longer horizon regressions.

2.1. *Data description and variance risk premium*

Our main data are from January 1990 to December 2010. We use a monthly frequency throughout this paper and thus have 252 observations available.

2.1.1. *Treasury data*

To consistently calculate short horizon excess returns on Treasuries, we use the Fama bond portfolios available from CRSP. The portfolios contain bonds

⁴Additional information on data construction is deferred to a separate appendix.

for issues maturing in a range around the quote dates. The portfolio returns are calculated as the equal-weighted average of the unadjusted holding period return for each bond in the portfolio. We calculate excess returns on a total of six portfolios with underlying maturities ranging from zero to 10 years.

Alternatively, we use the Fama T-bill structures from CRSP to compute short horizon excess returns on T-bills ranging between two and six months. To calculate one-year excess bond returns on longer maturity bonds we use the Fama and Bliss discount bond database from CRSP. We compute yields, returns, and forward rates for two to five year bonds. To have a consistent source of yields for calculating monthly and yearly excess returns, we also use the [Gürkaynak et al. \(2007, GSW dataset\)](#) dataset, which allows constructing one-month excess returns for longer maturity bonds.

Yields and returns are computed in logs. Yield spreads and excess returns are constructed relative to the one-period bond (one month for the portfolio and T-bill excess returns, one-year for Treasury bonds). We denote by $r_{t+1}^{(\tau)} = p_{t+1}^{(\tau-1)} - p_t^{(\tau)}$, the return on a τ year bond with log price $p_t^{(\tau)}$. The excess bond return is defined as

$$rx_{t+1}^{(\tau)} \equiv r_{t+1}^{(\tau)} - y_t^{(1)},$$

where $y_t^{(1)}$ is the one-period yield.

From the Fama and Bliss discount bond data, we also construct a tent-shaped factor from forward rates and the [Cochrane and Piazzesi \(2005\)](#) factor, CP. In order to construct the CP factor until December 2010 we need Fama Bliss Treasury bond data until December 2011. Overall, this restricts our sample to end in December 2010.

[Wright and Zhou \(2009\)](#) document the strong predictive power of the mean jump size for bond risk premia, and accordingly, we measure the 24-month rolling realized jump mean, \tilde{J} , using five-minute frequency data on the 30-year Treasury bond futures, under the assumption that jumps are rare and large.

2.1.2. *Implied variance data*

As has become standard practice, we use the squared VIX to proxy for the risk-neutral expectation of equity return variance for the next 30 days. The squared VIX is the model-free implied variance of the S&P 500 index calculated using S&P 500 index options.⁵ We use end-of-month data from the Chicago Board of Options Exchange (CBOE).

⁵See also [Demeterfi et al. \(1999\)](#); [Britten-Jones and Neuberger \(2000\)](#). The VIX White Paper (CBOE, 2003) outlines the calculation procedure.

2.1.3. *Stock index data*

To calculate the objective expectation we use intra-day data for the S&P 500 index sampled at the five-minute interval as in Bollerslev *et al.* (2009). The intra-day data are obtained from Tickdata.

2.1.4. *Macroeconomic data*

We compute the eight static macroeconomic factors $\widehat{F}_j, j = 1, \dots, 8$, following Ludvigson and Ng (2009, 2010). We update the time series and exclude the stock market and interest rate time series in order to have pure macro factors.⁶ The macroeconomic data are mainly from Global Insight.⁷

Summary statistics for the bond and bond portfolio returns are in Table 1, Panels A and B. Summary statistics for the macro control variables are collected in Panel C. The mean portfolio and bond returns are increasing with maturity and the numbers are in line with previous studies. While long-term bond excess returns — two- to five-year bonds for a one-year holding period — are highly persistent, possibly due to the overlapping return horizon of 11 months, the autocorrelation coefficients for the one-month holding period returns of the bond portfolios and the excess returns on Treasury bills are much lower with values of first-order autocorrelations, ranging from 0.11 to 0.75 for the bond portfolios and from 0.26 to 0.56 for T-bill excess returns.

Table 1. Summary statistics.

Panel A: Monthly Bond Portfolio Returns						
	<1y	1y–2y	2y–3y	3y–4y	4y–5y	5y–10y
Mean	0.35	0.42	0.48	0.52	0.55	0.57
Max	0.89	1.66	2.44	3.10	4.30	6.45
Min	-0.06	-0.72	-1.64	-2.47	-3.23	-3.58
StDev	0.21	0.42	0.70	0.96	1.21	1.47
Skewness	0.07	0.11	-0.11	-0.19	-0.18	-0.04
Kurtosis	2.40	3.15	3.19	3.14	3.31	3.80
AC(1)	0.75	0.34	0.24	0.18	0.16	0.11
AC(2)	0.74	0.25	0.10	0.03	-0.02	-0.09

⁶The original data set was previously used in Stock and Watson (2002). The stock market and interest rate time series we exclude are the Ludvigson and Ng (2009) series 82 through 102. In addition, we have to exclude seven variables that are no longer available after 2007. Consequently, we use 104 instead of 132 macroeconomic time series. For a shorter sample period ending in 2007 we use the original factors from Ludvigson and Ng (2010) as a robustness check. Our main results remain unchanged. We defer a more detailed description of the data to a separate appendix.

⁷In addition, three series are from the BEA and one is from the University of Michigan.

Table 1. (Continued)

Panel B: T-Bill (1 month) and Treasury Bond (1 year) Excess Returns									
	2m	3m	4m	5m	6m	2y	3y	4y	5y
Mean	0.30	0.42	0.41	0.63	0.71	0.95	1.80	2.54	2.96
Max	2.50	3.31	4.54	5.81	7.65	3.64	7.31	10.30	12.54
Min	-0.57	-0.87	-1.19	-1.43	-6.07	-2.37	-5.24	-6.88	-8.37
StDev	0.42	0.52	0.71	0.89	1.19	1.36	2.58	3.60	4.43
Skewness	1.83	1.66	1.83	1.78	0.87	-0.12	-0.27	-0.34	-0.42
Kurtosis	8.00	7.52	9.23	8.34	11.68	2.17	2.43	2.52	2.66
AC(1)	0.56	0.34	0.24	0.30	0.27	0.95	0.94	0.93	0.92
AC(2)	0.41	0.28	0.22	0.22	0.22	0.88	0.86	0.84	0.82

Panel C: Macro Variables										
	CP	\tilde{J}	\hat{F}_1	\hat{F}_2	\hat{F}_3	\hat{F}_4	\hat{F}_5	\hat{F}_6	\hat{F}_7	\hat{F}_8
Mean	1.02	0.01	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Max	5.49	0.13	22.54	14.71	9.17	10.18	9.14	5.54	5.70	6.43
Min	-4.27	-0.14	-7.94	-12.19	-6.21	-9.67	-14.36	-5.24	-5.09	-4.45
StDev	1.54	0.05	4.70	3.19	2.89	2.19	1.95	1.86	1.77	1.65
Skewness	0.34	-0.68	1.71	0.70	0.95	0.13	-0.73	0.19	-0.13	0.02
Kurtosis	3.33	3.31	7.42	6.93	4.26	5.94	17.43	3.36	3.41	3.46
AC(1)	0.92	0.92	0.85	-0.13	0.66	0.38	-0.28	0.45	0.15	-0.23
AC(2)	0.85	0.82	0.84	-0.28	0.72	0.46	-0.08	0.40	0.24	0.12

Panel D: Variance Risk Premia					
	VIX	RV	Proj	VRP(RV)	VRP(Proj)
Mean	20.41	14.21	14.10	17.15	21.57
Max	59.89	73.11	35.36	116.85	194.72
Min	10.42	4.43	9.22	-298.37	0.18
StDev	7.88	8.31	4.54	18.31	28.64
Skewness	1.57	2.68	2.16	-5.36	3.27
Kurtosis	6.91	14.96	8.66	62.00	19.22
AC(1)	0.86	0.80	0.82	0.33	0.76
AC(2)	0.72	0.62	0.68	-0.05	0.56

Notes: This table presents summary statistics for all monthly data from January 1990 to December 2010. In Panels A and B, we report the one-month holding period returns on bond portfolios, one-month excess returns on Treasury bills, and one-year excess returns on Treasury Bonds. Panel C reports summary statistics for the CP factor, CP, the mean jump size, \tilde{J} , and the eight Ludvigson and Ng macro factors, \hat{F}_j . Panel D reports the summary statistics for the VIX, the expected realized volatility and the actual realized volatility. Values are expressed in percentages and annualized. The variance risk premia are calculated as the difference between the square of the monthly implied and realized volatilities expressed in percent.

Table 2. Cross-correlations of predictor variables.

	VRP	CP	\tilde{j}
VRP	1		
CP	-0.225	1	
\tilde{j}	-0.061	-0.235	1
\widehat{F}_1	0.470	-0.031	-0.283
\widehat{F}_2	0.161	-0.010	-0.085
\widehat{F}_3	0.045	0.140	-0.140
\widehat{F}_4	0.154	0.230	0.008
\widehat{F}_5	0.038	0.088	-0.038
\widehat{F}_6	-0.097	0.082	-0.039
\widehat{F}_7	-0.001	0.259	0.047
\widehat{F}_8	0.090	-0.167	0.063

Notes: This table presents the cross-correlation for the CP factor, CP, the mean jump size, \tilde{j} , the eight Ludvigson and Ng macro factors, \widehat{F}_j , and the variance risk premium, VRP. We use monthly data from January 1990 to December 2012.

The CP factor and the mean jump size are highly persistent with first-order autocorrelation coefficients of 0.92. The macro factors, \widehat{F}_j display much lower autocorrelations on average and some factors even display a negative autocorrelation.

Table 2 reports the unconditional correlation among all the predictor variables, including the variance risk premium. The macro factors are calculated specifically for the 1990 to 2010 period and hence, they all have zero cross-correlations. The variance risk premium is not very highly correlated with the other factors except for the CP factor (-0.22), and the first macro component, \widehat{F}_1 (0.47), which Ludvigson and Ng (2009) label as the ‘real factor’ due to its high correlation with measures of real output and employment. This echoes the finding in Bollerslev and Zhou (2007) that the variance risk premium may be intimately related to economic fundamentals in terms of the uncertainty shocks.

2.1.5. Forecasting realized variance and the variance risk premium

To estimate realized variance, we use high frequency data for the S&P 500 index as the VIX — our measure of implied volatility — is calculated using options on the S&P 500 index. Let $RV_{t,\tau}$ be the realized variance from day $t - \tau$ to day t , with τ being typically a month or the equivalent of 21 trading days. To estimate the objective expectation of return variation of the next

period, $\mathbb{E}_t^{\mathbb{P}}(\text{RV}_{t+\tau,\tau})$, we first consider the realized variance RV_t at day t , which is defined as

$$\text{RV}_t = \sum_{i=1}^M r_{t,i}^2, \tag{1}$$

where $r_{t,i} = \log P(t - 1 + \frac{i}{M}) - \log P(t - 1 + \frac{i-1}{M})$ is the intra-daily log return in the i th sub-interval of day t and $P(t - 1 + i/M)$ is the asset price at time $t - 1 + i/M$. For each day, we take $r_{t,i}$ between 9:00 and 15:00 at every five-minute interval to calculate RV_t . In addition, we also include the overnight return in the calculation of the realized variance. The normalized monthly realized variation, $\text{RV}_{t,\text{mon}}$, is defined by the average of the 21 daily measurements $\text{RV}_{t,\text{mon}} = \frac{1}{21} \sum_{j=0}^{20} \text{RV}_{t-j}$. The normalized weekly realized variation $\text{RV}_{t,\text{week}}$ is correspondingly defined by the average of the five daily measurements $\text{RV}_{t,\text{week}} = \frac{1}{5} \sum_{j=0}^4 \text{RV}_{t-j}$.

To better capture the long memory behavior of volatility, we use the daily, weekly and monthly realized variance estimates to estimate the heterogeneous autoregressive model of realized volatility (HAR-RV) proposed by Corsi (2009). HAR-RV estimators have become increasingly popular in financial econometrics literature in the past years (see Corsi *et al.* (2010), Bollerslev and Todorov (2011) and Patton and Sheppard (2011)). HAR-RV is a parsimonious version of high-order auto-regressions. We augment the HAR monthly forecasting model with additional lags of implied variance

$$\begin{aligned} \text{RV}_{t+21,\text{mon}} = & \alpha + \beta_D \text{RV}_t + \beta_W \text{RV}_{t,\text{week}} + \beta_M \text{RV}_{t,\text{mon}} + \sum_{i=1}^k \beta_{V,i} \text{VIX}_{t-i}^2 \\ & + \epsilon_{t+21,\text{mon}}, \end{aligned} \tag{2}$$

where VIX_t^2 is the square of the daily VIX index divided by $12 \times 10^4 \times 30$ to be comparable to $\text{RV}_{t,\text{mon}}$. Equation (2) is motivated by the large literature in derivatives pricing, showing that implied variance is a more efficient forecast for future realized variance than its own lag (Jiang and Tian, 2005) and extends the forecasting model of Drechsler and Yaron (2011) that uses one lag realized variance and one lag implied variance.

In our implementation of the HAR-RV model we are careful to ensure that the forecast of monthly realized variance, denoted RV_t^{HAR} , can be obtained in real time and does not suffer from any look ahead bias. Thus, we implement the regression 2 using an expanding sample of data to obtain a true out-of-sample forecast of the future month's realized variance. This requires a burn-in period to first run the regression. In order to limit the loss

of available data and because VIX is only available starting in January 1990, we use the squared VXO as our implied variance proxy on the RHS of the regression. Unlike the VIX, the VXO is not calculated using a model-free approach but is simply the Black–Scholes implied volatility of at-the-money S&P 100 options. However, the correlation between VIX and VXO is almost perfect and, in sample regressions for the common sample period between 1990 and 2010, lead to essentially the same results whether we use the VIX or the VXO.⁸

The variance risk premium is formally defined as the difference between the expected future variation under the risk-neutral and actual probability measures between day t and T

$$\text{VRP}_{t,\tau} \equiv \mathbb{E}_t^{\mathbb{Q}}(\text{RV}_{t+\tau,\tau}) - \mathbb{E}_t^{\mathbb{P}}(\text{RV}_{t+\tau,\tau}), \quad (3)$$

where $\tau = T - t$ denotes the time horizon which typically is a month or 21 trading days.⁹

As discussed above, we use the expanding projection to proxy for the expected realized variance. Hence, $\mathbb{E}_t^{\mathbb{P}}(\text{RV}_{t+\tau,\tau}) = \text{RV}_t^{\text{HAR}}$. To proxy for the risk-neutral variance, we take the VIX squared of the S&P 500 index with a one-month horizon, using a model-free approach. With some regularity assumptions and even if the underlying asset follows a general jump diffusion (see [Jiang and Tian, 2005](#); [Carr and Wu, 2009](#)), this risk-neutral expected variance can be computed as a portfolio of European calls on the underlying.

We plot the VIX, the expected realized volatility (the square root of RV_t^{HAR}), together with the variance risk premium in [Fig. 1](#). Summary statistics are reported in [Table 1](#), Panel D. The VIX and the expected realized volatility are expressed in percent and annualized. In addition, we also report the summary statistics for the actual realized volatility, also expressed in percent and annualized. The variance risk premium is obtained by taking the differences of the squared monthly implied and realized variances expressed in percent.¹⁰ We report summary statistics for the *ex ante* variance risk premium calculated using our projection and the *ex post* realized variance risk

⁸Adding the implied variance data to the forecasting regression leads to marginally better forecasts. However, an implementation, using only the standard HAR-RV model, leads to a very similar time series of expected realized variances.

⁹For notational simplicity, we subsequently drop the subscript τ as we always consider the one-month horizon variance risk premium VRP_t .

¹⁰Or equivalently, it is the monthly variance risk premia expressed in decimals and then multiplied by 10^4 .




Fig. 1. Market variance risk premium.

Notes: The upper panel plots the VIX together with the forecasted realized volatility (i.e., the square root of the expected realized variance) which we calculate from a projection for the realized variance:

$$RV_{t+21,mon} = \alpha + \beta_D RV_t + \beta_W RV_{t,week} + \beta_M RV_{t,mon} + \sum_{i=1}^k \beta_{V,i} VIX_t^2 + \epsilon_{t+21,mon},$$

where $RV_{t,week} = 1/5 \sum_{j=0}^4 RV_{t-j}$, $RV_{t,mon} = 1/21 \sum_{j=0}^{20} RV_{t-j}$ and VIX_t^2 is the square of the daily VIX index divided by 12×10^4 to convert numbers into a monthly quantity that is comparable to $RV_{t,mon}$. RV_t represents the daily realized variance calculated using 5-min squared returns on the S&P 500 index.

premium calculated using the future month realized variance instead of the forecast.

The figure reveals that most of the peaks in the variance risk premium occur during periods of financial crises such as the LTCM default in August 1998, the burst of the dot com bubble in 2000, and the most recent financial crisis in late 2008.

Looking at the summary statistics in Table 1, Panel D, there are several interesting points to highlight. First, the average realized variance risk premium (last column, VRP^{HAR}) is around 22% which is comparable to the numbers found previously in the literature (see, e.g., Drechsler and Yaron, 2011; Bekaert and Engstrom, 2009; Bekaert *et al.*, 2010). The realized variance risk premium VRP^{5min} is slightly lower at about 17%. Second, contrary to the previous literature, our measure never turns negative, which is partly driven by our HAR-RV forecasting model, which is augmented by 4 lags of implied variance. This is important not only from an empirical point of view but also given the theoretical underpinnings — the variance risk premium is

usually interpreted as an insurance premium for investors who pay for an asset whose payoff is high when return variation is large. Not surprisingly, the realized variance risk premium occasionally turns very negative. It also exhibits very negative skewness and extremely high kurtosis. Third, the first and second-order autocorrelation coefficients of our variance risk premium are 0.76 and 0.56, which alleviates econometric concerns of regressing on highly persistent variables. The realized variance risk premium is even less persistent with autocorrelation coefficients of only 0.33 and -0.05 , respectively.

2.2. Short and long horizon bond return predictability

Next, we document the predictive power of the equity variance risk premium for bond excess returns. First, we show the predictability for short horizon bond excess returns regardless of the underlying maturity of the bonds. Then, we show that while the variance risk premium may be able to predict short-run returns it largely fails to predict longer horizon bond returns.

We start by running the following regressions for Treasury bond portfolios

$$rx_{t+h}^{(\tau)} = \beta_0^{(\tau)}(h) + \beta_1^{(\tau)}(h)VRP_t + \epsilon_{t+h}^{(\tau)}, \quad (4)$$

where $rx^{(\tau)}$ is the one month excess return ($h = 1$) on one of the six bond portfolios with underlying maturities $\tau = < 1y, 1y - 2y, 2y - 3y, 3y - 4y, 4y - 5y, 5y - 10y$ and VRP_t is the equity variance risk premium. The estimated slope coefficients from the regression are summarized in the upper left panel of Fig. 2. We plot standardized coefficients, meaning that all variables have zero mean and a standard deviation of one to make coefficients comparable in terms of the economic significance as well as the statistical significance, summarized by the 95% confidence band. The standardized coefficients range between 0.15 and 0.25 and all are highly significant. We repeat the same regressions using one-month excess returns on $\tau = 2m, 3m, 6m, 1y, 2y, 3y, 4y, 5y, 10y$ bonds calculated using the [Gürkaynak et al. \(2007\)](#) dataset (GSW dataset henceforth). The estimated standardized coefficients are again highly significant and range between 0.13 and 0.18. The results are plotted in the upper right panel of Fig. 2.

We repeat the regressions using a longer holding period of one year for the same underlying assets, i.e., we first regress one-year excess returns on the bond portfolios on the variance risk premium and then we repeat the regressions for one-year bond excess returns calculated using the GSW dataset. The results are plotted in the lower left and right panels of Fig. 2. Unlike for short horizon regressions, the variance risk premium is not a good

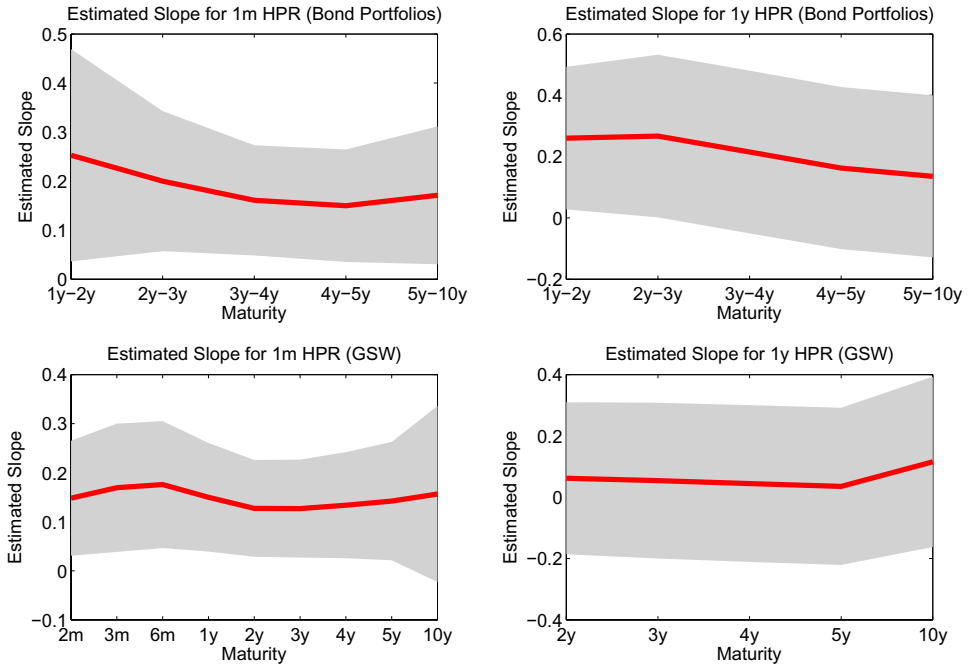


Fig. 2. Estimated slope coefficients from univariate regressions.

Notes: The upper left panel plots the estimated slope coefficient from regressing monthly excess returns of bond portfolios on the variance risk premium. The upper right panel plots the estimated slope coefficient from the same regression for one-year bond portfolio excess returns. The lower two panels plot the coefficients for the same univariate regressions using one-month and one-year excess returns for Treasury bonds calculated using the [Gürkaynak et al. \(2007\)](#) dataset. The coefficient estimates are from standardized regressions (all variables have zero mean and a standard deviation of one) to visualize the economic significance. The shaded areas represent the 95% confidence bounds. Coefficients are estimated using monthly data from January 1990 to December 2010.

predictor of longer horizon excess returns. The standardized coefficients for the GSW bond excess returns are very small and not significant. For the bond portfolios, the variance risk premium has marginal predictive power for portfolios with short underlying maturities, while no predictability exists for longer maturity portfolios.

In essence, we find that the variance risk premium predicts bond risk premia at short horizons well, but has little predictive power at longer horizons. The R^2 for the univariate regressions for short horizons range between 2% and 6% for the bond portfolios and between 2% and 3% for the GSW excess returns. The predictability seems independent of the underlying maturity of the bonds and is only a function of the horizon.

To further test this preliminary empirical regularity, we turn to two sets of Treasury data that have been extensively studied in the literature, Treasury

bills and bonds. We repeat regression 4 for the Treasury bill excess returns for maturities $\tau = 2, 3, 4, 5$ and 6 months. We mainly focus on the one-month horizon regressions but we also consider the additional horizons $h = 2, 3, 4$ and 5 months. Furthermore, we run regressions for one-year excess returns ($h = 12$) on Fama Bliss discount bonds with maturities $\tau = 2, 3, 4$ and 5 years.

Given the data available from CRSP, it is not possible to construct one-month excess returns on the Fama Bliss bonds. Thus, it is not possible to test whether the variance risk premium predicts short horizon excess returns on Fama Bliss Treasury bonds. Similarly, there are some shortcomings to using Treasury bills as well. Duffee (1996) documents a dramatic weakening of the links between Treasury bill yields and yields on other Treasury securities. There is significant idiosyncratic behavior in the shortest maturity yields, which may be partly due to increased market segmentation. Duffee for example argues that term structure models should not be calibrated using one-month Treasury bill yields. This suggest that the results from the Treasury bill regression may have to be treated with some caution.

When running the Treasury bill regression for the full sample period of 1990 to 2010, we indeed found that the variance risk premium is not a significant predictor for the shortest maturity bills and the variance risk premium is only significant for four-month maturity Treasury bills and beyond. The breakdown of the predictive relationship between the variance risk premium and Treasury bill excess returns can be attributed to the crisis and is possibly a result of the significant activism of the Fed after the Lehman bankruptcy in September 2008 and its effect on short-term yields.

In Table 3, Panel A we thus present regression results for the Treasury bill regressions, for a sample period that ends with the Lehman bankruptcy in September 2008. The coefficients in the table are not standardized but they are obtained from regressing excess returns on the variance risk premium and a constant. However, we only present the coefficient estimates and the adjusted R^2 in the table. The coefficient estimates are all significant for one-month excess returns while the variance risk premium loses its predictive power when moving to horizons of two months and above. In terms of economic significance, the standardized coefficients for the one-month excess return regressions decrease with maturity of the underlying Treasury bills and range between 0.17 and 0.31.

Panel A also contains the coefficient estimates for the one-year excess return regressions using Fama Bliss discount bonds. Here, the estimates do not depend on the sample period and we present the full sample regressions

Table 3. T-bill and bond regressions.

Panel A: Univariate Regressions						
h	T-Bills	2m	3m	4m	5m	6m
1m	VRP	0.078 (2.51)	0.065 (2.30)	0.070 (2.05)	0.093 (2.20)	0.125 (2.18)
	Adj. R^2	0.09	0.05	0.03	0.03	0.03
2m	VRP		0.040 (1.43)	0.046 (1.38)	0.056 (1.01)	0.106 (1.26)
	Adj. R^2		0.05	0.01	0.01	0.01
3m	VRP			0.031 (0.93)	0.042 (0.73)	0.084 (0.90)
	Adj. R^2			0.01	0.00	0.01
4m	VRP				0.034 (0.69)	0.073 (0.86)
	Adj. R^2				0.00	0.01
5m	VRP					0.054 (0.92)
	Adj. R^2					0.02
1y	Treasury bonds		2y	3y	4y	5y
	VRP		0.033 (0.44)	0.072 (0.51)	0.040 (0.20)	0.052 (0.21)
	Adj. R^2		-0.00	0.00	-0.00	-0.00

Panel B: Multivariate Regressions						
	T-Bills	2m	3m	4m	5m	6m
1m	VRP	0.079 (3.87)	0.058 (3.15)	0.058 (2.12)	0.080 (2.65)	0.105 (2.87)
	Adj. R^2	0.32	0.23	0.23	0.25	0.26
1y	Treasury bonds		2y	3y	4y	5y
	VRP		-0.015 (-0.25)	0.016 (0.13)	0.029 (0.17)	0.081 (0.36)
	Adj. R^2		0.42	0.42	0.43	0.42

Notes: For Panel A, we run the following regression: $rx_{t+h}^{(\tau)} = \beta_0^{(\tau)}(h) + \beta_1^{(\tau)}(h)VRP_t + \epsilon_{t+h}^{(\tau)}$, where $rx^{(\tau)(h)}$ are either the excess returns on Treasury bills, $h = 1, 2, 3, 4, 5$ and $\tau = 2, 3, 4, 5, 6$ months, or the excess returns on Treasury bonds, $h = 12$ and $\tau = 24, 36, 48, 60$ months. VRP_t is the market variance risk premium. For Panel B, the regression is augmented with the CP factor, the jump factor and the eight Ludvigson and Ng macro factors. The table only reports the coefficient estimate for the variance risk premium and the overall adjusted R^2 . Coefficients are estimated with ordinary-least squares. Standard errors are in parentheses and are calculated using Newey and West (1987) standard errors. The sample spans the period from January 1990 to December 2010.

until December 2010. The results mirror those obtained with the GSW data: none of the coefficients are either economically or statistically significant and R^2 are essentially zero.

To summarize, the results using Treasury bill and Fama Bliss Treasury bond data are largely consistent with the overall picture, namely that the equity variance risk premium predicts short-run bond excess returns. In contrast, the variance risk premium has little predictive power for longer horizon excess returns. These findings seem to be complementary to a large body of existing literature because the market variance risk premium may capture a unique component of bond risk premia that is relevant for the short horizon and is driven by economic uncertainty shocks, but that is at the same time orthogonal to the long horizon component captured by established predictors. To further justify our empirical conjecture, we next turn to a host of robustness checks involving well established bond return predictors.

2.3. Controlling for other bond return predictors

A large literature has been devoted to studying different factors that predict bond risk premia at long horizons (usually one year) and it is thus, natural to question whether the predictive power of the variance risk premium is also subsumed by those predictor variables. Our analysis is focused on three sets of additional predictor variables: The CP factor from [Cochrane and Piazzesi \(2005\)](#), the jump factor from [Wright and Zhou \(2009\)](#) and the macro factors from [Ludvigson and Ng \(2009\)](#). [Cochrane and Piazzesi \(2005\)](#) find that a linear combination of forward rates is the most powerful predictor for long-term bond returns. [Wright and Zhou \(2009\)](#) show that the mean jump size explains a significant fraction of the variation in long-term bond excess returns and doubles the adjusted R^2 when combined with the CP factor. [Ludvigson and Ng \(2009\)](#) include macro factors extracted from a large set of macro variables, using principal components analysis to explain a highly significant fraction of the time variation in bond excess returns. Additionally, [Duffee \(2011\)](#) estimates a latent factor from a five factor Gaussian model which has predictive power for bond excess returns but is not (or only weakly) spanned by the cross-section of yields. [Cieslak and Povala \(2010\)](#) find high R^2 s when running predictive regressions from bond excess returns on cycles which represent deviations from the long-run relationship between yields and the slow-moving component of inflation and savings.¹¹ Our extended

¹¹For a shorter sample, we can also include the hidden factor in our analysis. The results remain unchanged with respect to the variance risk premium.

regression is of the following form:

$$rx_{t+h}^{(\tau)} = \beta_0^{(\tau)}(h) + \beta_1^{(\tau)}(h)\text{VRP}_t + \beta_2^{(\tau)}(h)\text{CP}_t + \beta_3^{(\tau)}(h)\tilde{J}_t + \sum_{j=1}^8 \beta_{3+j}^{(\tau)}(h)\hat{F}_{j,t} + \epsilon_{t+h}^{(\tau)},$$

which is simply the univariate regression 4 augmented by the Cochrane–Piazzesi factor, CP, the mean jump size, \tilde{J} , and the Ludvigson and Ng macro factors, \hat{F}_j .

We present the multivariate regression results for bond portfolio excess returns in Table 4. The main result is that the variance risk premium is robust to including a host of additional predictors in the multivariate regressions. Statistical significance is even slightly higher than for the univariate regressions and the economic significance remains virtually unchanged with standardized coefficients ranging between 0.16 and 0.31. Thus, for the one-month holding period, adding other predictors does not change the significance of the variance risk premium very much. The striking short-run predictability of the variance risk premium for bond returns is similar to the one reported by Zhou (2010) and mirrors the findings for stock returns by Bollerslev *et al.* (2009) and Drechsler and Yaron (2011).

As for the other predictors in bond portfolio regressions, the CP also shows up as a significant predictor with statistical and economic significance in roughly the same order of magnitude as the variance risk premium. In addition, the jump factor is marginally significant for excess returns on some shorter maturity portfolios but loses its significance for longer maturities. Both have the same sign as the known regressions for one-year bond excess returns. In addition, some of the LN factors are occasionally significant although there does not seem to exist a consistent pattern, except for \hat{F}_2 which is significant throughout. Adding the additional regressors increases adjusted R^2 to between 10% and up to almost 40%.

The same pattern emerges when running multivariate regressions on GSW short-run bond excess returns.¹² The variance risk premia remains highly significant for bond excess returns for all underlying maturities. The CP factor picks up some predictability and the jump factor is significant for some intermediate maturities. R^2 are slightly lower and range between 7% and 18%.

Not very surprisingly, the variance risk premium does not regain predictive power for one-year excess returns when adding the additional control

¹²These results are not reported.

Table 4. Multivariate excess bond portfolio returns — one month.

	<1y	1y-2y	2y-3y	3y-4y	4y-5y	5y-10y	
Constant	-3.951 (-9.90)	-4.111 (-7.41)	-4.174 (-4.89)	-4.486 (-3.80)	-5.133 (-3.38)	-5.904 (-2.96)	-6.648 (-3.07)
VRP	0.292 (2.69)	0.536 (3.43)	0.737 (3.48)	0.963 (3.08)	1.313 (2.71)	1.760 (2.47)	1.908 (2.72)
CP	0.335 (2.52)	0.344 (2.98)	0.795 (3.44)	1.239 (3.14)	1.780 (3.12)	2.145 (2.96)	2.441 (2.87)
\tilde{J}	0.382 (0.09)	-11.844 (-2.07)	-17.794 (-1.94)	-20.339 (-1.52)	-21.675 (-1.21)	-18.684 (-0.77)	-18.684 (-0.77)
\hat{F}_1	0.145 (3.45)	0.199 (3.19)	0.214 (2.10)	0.147 (1.02)	0.047 (0.23)	-0.181 (-0.62)	-0.181 (-0.62)
\hat{F}_2	0.073 (1.94)	0.201 (2.61)	0.293 (2.49)	0.423 (2.81)	0.604 (2.96)	0.855 (2.60)	0.855 (2.60)
\hat{F}_3	0.169 (3.26)	0.149 (2.70)	0.154 (1.58)	0.201 (1.36)	0.237 (1.21)	0.430 (1.57)	0.430 (1.57)
\hat{F}_4	-0.196 (-3.18)	-0.330 (-3.11)	-0.605 (-3.11)	-0.869 (-3.07)	-1.006 (-2.60)	-0.852 (-1.35)	-0.852 (-1.35)
\hat{F}_5	0.058 (1.36)	0.221 (1.73)	0.349 (1.33)	0.520 (1.44)	0.417 (0.80)	-0.033 (-0.04)	-0.033 (-0.04)
\hat{F}_6	0.248 (2.82)	0.215 (1.44)	0.190 (0.81)	0.144 (0.49)	0.091 (0.25)	-0.159 (-0.38)	-0.159 (-0.38)

Table 4. (Continued)

	<1y	1y-2y	2y-3y	3y-4y	4y-5y	5y-10y
\widehat{F}_7	-0.255 (-3.79)	-0.402 (-3.03)	-0.667 (-2.77)	-0.934 (-2.73)	-1.300 (-2.85)	-1.693 (-2.72)
\widehat{F}_8	-0.221 (-3.20)	-0.262 (-1.84)	-0.324 (-1.47)	-0.354 (-1.22)	-0.352 (-0.95)	-0.277 (-0.60)
R^2	0.114	0.398	0.070	0.180	0.062	0.139
Adj. R^2	0.107	0.092	0.063	0.143	0.054	0.057

Notes: We run the following regression: $rx_{t+1} = \beta_0^{(\tau)} + \beta_1^{(\tau)}(h)VRP_t + \beta_2^{(\tau)}CP_t + \beta_3^{(\tau)}\bar{J}_t + \sum_{j=1}^8 \beta_{3+j}^{(\tau)}\widehat{F}_{j,t} + \epsilon_{t+1}^{(\tau)}$, where $rx^{(\tau)}$ is the one-month excess return ($h = 1$) on one of the six bond portfolios with underlying maturities, $\tau = < 1y, 1y - 2y, 2y - 3y, 3y - 4y, 4y - 5y, 5y - 10y$, and VRP_t is the equity variance risk premium. CP is the Cochrane-Piazzesi factor, \bar{J} is the mean jump size, and \widehat{F}_j denotes the eight Ludvigson and Ng macro factors. Coefficients are estimated with ordinary-least squares. Standard errors are in parentheses and are calculated using Newey and West (1987) standard errors. The sample spans the period from January 1990 to December 2010.

variables and the results summarized in Fig. 2 remain unchanged. Thus, we do not report the detailed multivariate regression results for one-year excess returns here. To summarize, for both the portfolio and the GSW bond excess returns, adding the CP factor, the jump size, and the macro variables to the regression raises R^2 to between 45% and almost 60%. The CP and the jump factor become highly significant for the GSW bond excess returns, while most of the increase in R^2 for bond portfolio excess returns can be attributed to the inclusion of the macro variables.

We also include the additional predictor variables in the Treasury bill and Fama Bliss Treasury bond regressions. The relevant results are summarized in Panel B of Table 3. Again, we focus on the estimated coefficients for the variance risk premium only. We omit the results for the Treasury bill regressions for holding periods of two months and beyond as the variance risk premium is not significant in univariate regressions to begin with. However, the variance risk premium remains highly significant for Treasury bill regressions at the one-month horizon. The economic significance remains virtually unchanged and standardized coefficient estimates range between 0.14 and 0.32. Adding all the additional predictor variables significantly raises R^2 but the CP factor is no longer significant once the jump factor and the macro variables are added to the regression. The jump factor is increasingly significant with increasing maturity of the Treasury bills and various macro factors are consistently significant (yet notoriously hard to interpret).

The results for the Fama Bliss Treasury bond regressions are not very surprising at this point. The variance risk premium has zero predictive power while the R^2 are increased to more than 40% once the additional predictors are included in the analysis. The CP factor and the jump mean are both significant and capture most of the predictability. Once the jump mean is included in the regression, adding the macro variables does not add much to the picture.

Overall, a fairly consistent pattern emerges. The variance risk premium contains relevant information for short-run bond risk premia for bonds of any maturity, while it has no predictive power for longer horizon excess returns. The contrast between short and long-term bond risk premia is best seen in Fig. 3. It is clear that short-term bond risk premia (top panel) have large spikes around major financial crises and economic recessions, but these shocks are generally short-lived — uncertainty comes and goes. On the other hand, the long-term bond risk premia (bottom panel) seem to have gradual persistent swings at least several years apart and sometimes even as long as the




Fig. 3. *Ex post* versus *ex ante* bond risk premia in short and long horizons.

Notes: The upper panel plots the average one-month bond risk premium for Treasury bills with maturities of two to six months (dashed line) together with the fitted value from a regression (bold line). The lower panel plots the average one-year bond risk premium for Treasury bonds with maturities of two to five years (dashed line) together with the fitted value from a regression (bold line).

business cycle frequency — like the 2001–2008 cycle. Put together, we have a whole picture of bond risk premia responding both slowly to long-term cyclical risk and quickly to short-term uncertainty shocks.

In summary, we find that the stock market variance risk premium is a robust predictor of bond returns at the short end, but the predictive power becomes weaker at longer horizons, as also previously reported in [Baele et al. \(2010\)](#). The short-term predictive power of the variance risk premium is also robust to the inclusion of other standard predictors such as the CP factor, the jump mean, and macro variables. This is even more notable as the variance risk premium we calculate can be obtained in real-time and does not contain forward-looking information, whereas the other factors we use are extracted using information from the whole sample period.

3. Economic Uncertainty and Inflation Dynamics

To understand why the variance risk premium has significant predictive power for short-run bond risk premia, we present a stylized structural model where the real bond risk premium is present because of agents' preference for an early resolution of uncertainty. The nominal bond risk premia is non-

redundant because the inflation process co-varies with both cash flow and uncertainty shocks. The nominal bond risk premium in our economy works only through the conditional volatility channel and does not rely on the conditional mean channel (as in Pennacchi, 1991; Sun, 1992). As such, our model may be viewed as an extension of the consumption uncertainty model by Bollerslev *et al.* (2009).

Our calibration result suggests that the proposed inflation uncertainty model not only has the capability to replicate the predictability pattern of the variance risk premium for bond risk premia documented in recent research, but also matches the level of bond risk premia that is typically hard to pin down in structural economic models. The volatility or uncertainty channel to resolve the ‘expectations hypothesis’ puzzles is in contrast with those relying on the consumption growth channel — e.g., habit formation (Wachter, 2006), long-run risk (Bansal and Shaliastovich, 2009), or rare disasters (Gabaix, 2009).

3.1. Economic uncertainty and variance risk premia

The representative agent in the economy has a Epstein–Zin–Weil recursive preference and has the value function V_t of life-time utility given as

$$V_t = [(1 - \delta)C_t^{\frac{1-\gamma}{\theta}} + \delta(E_t[V_{t+1}^{1-\gamma}])]^{\frac{\theta}{1-\gamma}}, \quad (5)$$

where C_t is consumption at time t , δ denotes the subjective discount factor, γ refers to the coefficient of risk aversion, $\theta = \frac{1-\gamma}{1-\frac{1}{\psi}}$, and ψ equals the intertemporal elasticity of substitution (IES). The key assumption that $\psi > 1$ hence $\theta < 0$ implies that agents prefer an earlier resolution of economic uncertainty, such that the uncertainty or volatility risk in asset markets carries a *positive* risk premium.

The log consumption growth and its volatility follow the joint dynamics

$$g_{t+1} = \mu_g + \sigma_{g,t}z_{g,t+1}, \quad (6)$$

$$\sigma_{g,t+1}^2 = a_\sigma + \rho_\sigma\sigma_{g,t}^2 + \sqrt{q_t}z_{\sigma,t+1}, \quad (7)$$

$$q_{t+1} = a_q + \rho_qq_t + \varphi_q\sqrt{q_t}z_{q,t+1}, \quad (8)$$

where $\mu_g > 0$ denotes the constant mean growth rate.¹³ The time-variation in $\sigma_{g,t+1}^2$ is one of the two components that drives the equity risk premium, or the “consumption risk”; while the time-variation in q_t is not only responsible

¹³The parameters satisfy $a_\sigma > 0$, $a_q > 0$, $|\rho_\sigma| < 1$, $|\rho_q| < 1$, $\varphi_q > 0$; $\{z_{g,t}\}$, $\{z_{\sigma,t}\}$, and $\{z_{q,t}\}$ are iid $\mathcal{N}(0, 1)$ processes, jointly independent of each other.

for the “uncertainty risk” component in equity risk premium, but also constitutes the main driver of variance and bond risk premia, as explained below.

Let w_t denote the logarithm of the wealth–consumption ratio of the asset that pays the consumption endowment, $\{C_{t+i}\}_{i=1}^\infty$; and conjecture a solution for w_t as an affine function of the state variables $\sigma_{g,t}^2$ and q_t . One can solve for the coefficients A_0 , A_σ and A_q using the standard [Campbell and Shiller \(1988\)](#) approximation $r_{t+1} = \kappa_0 + \kappa_1 w_{t+1} - w_t + g_{t+1}$. The restriction that $\psi > 1$ hence $\theta < 0$ implies that the impact coefficient associated with both consumption and volatility state variables is negative; i.e., $A_\sigma < 0$ and $A_q < 0$. So if consumption and uncertainty risks are high, the price–dividend ratio is low, hence, the risk premia are high. In response to high economic uncertainty risks, agents sell risky assets, and consequently the wealth–consumption ratio falls; so that risk premia rise.

The conditional variance of the time t to $t + 1$ return, $\sigma_{r,t}^2 \equiv \text{Var}_t(r_{t+1})$, is given by: $\sigma_{r,t}^2 = \sigma_{g,t}^2 + \kappa_1^2 (A_\sigma^2 + A_q^2 \varphi_q^2) q_t$. The variance risk premium can then be defined as the difference between risk-neutral and objective expectations of the return variance¹⁴

$$\begin{aligned} \text{VRP}_t &\equiv E_t^Q(\sigma_{r,t+1}^2) - E_t^P(\sigma_{r,t+1}^2) \\ &\approx (\theta - 1)\kappa_1 [A_\sigma + A_q \kappa_1^2 (A_\sigma^2 + A_q^2 \varphi_q^2) \varphi_q^2] q_t > 0. \end{aligned} \quad (9)$$

One key observation here is that any temporal variation in the endogenously generated variance risk premium is due solely to the volatility-of-volatility or economic uncertainty risk, q_t , but not the consumption growth risk, $\sigma_{g,t+1}^2$. Moreover, provided that $\theta < 0$, $A_\sigma < 0$, and $A_q < 0$, as would be implied by the agents’ preference of an earlier resolution of economic uncertainty, this difference between the risk-neutral and objective expectations of return variances is guaranteed to be positive. If consumption volatility is not stochastic or there is no recursive preference, the variance risk premium is zero by construction.

3.2. Inflation dynamics and bond return predictability

In order for the real economy model outlined above to have realistic implications for nominal bond risk premia, one needs to impose rich inflation dynamics, which are capable of incorporating stochastic volatility, money non-neutrality, and perhaps both cash flow and uncertainty shocks. Our

¹⁴The approximation comes from the fact that the model-implied risk-neutral conditional expectation cannot be computed in closed form, and a log-linear approximation is applied.

preferred specification for expected inflation π_t is

$$\pi_{t+1} = a_\pi + \rho_\pi \pi_t + \varphi_\pi z_{\pi,t+1} + \varphi_{\pi g} \sigma_{g,t} z_{g,t+1} + \varphi_{\pi \sigma} \sqrt{q_t} z_{\sigma,t+1}, \quad (10)$$

where ρ_π is the persistence and $\frac{a_\pi}{1-\rho_\pi}$ is the long-run level of the inflation process. The innovations in the inflation dynamics consist of three parts: (1) a constant volatility part φ_π with exogenous shock $z_{\pi,t+1}$ that is uncorrelated with all shocks in the real model, (2) a stochastic volatility part $\varphi_{\pi g} \sigma_{g,t}$ that works through the consumption growth channel $z_{g,t+1}$, and (3) another stochastic volatility part $\varphi_{\pi \sigma} \sqrt{q_t}$ that works through the volatility channel $z_{\sigma,t+1}$. Note that $\varphi_{\pi g}$ and $\varphi_{\pi \sigma}$ “leverage up” the inflation exposure to the growth and uncertainty risks. Hence, money-neutrality is implicitly violated.

We can examine each component of inflation shocks separately to assess which channel affects the bond risk premia more and to what degree it does so.

$$\begin{aligned} \text{Model I : } \pi_{t+1} &= a_\pi + \rho_\pi \pi_t + \varphi_\pi z_{\pi,t+1}, \\ \text{Model II : } \pi_{t+1} &= a_\pi + \rho_\pi \pi_t + \varphi_\pi \sigma_{g,t} z_{\pi,t+1}, \\ \text{Model III : } \pi_{t+1} &= a_\pi + \rho_\pi \pi_t + \varphi_\pi z_{\pi,t+1} + \varphi_{\pi g} \sigma_{g,t} z_{g,t+1}, \\ \text{Model IV : } \pi_{t+1} &= a_\pi + \rho_\pi \pi_t + \varphi_\pi z_{\pi,t+1} + \varphi_{\pi \sigma} \sqrt{q_t} z_{\sigma,t+1}, \\ \text{Model V : } \pi_{t+1} &= a_\pi + \rho_\pi \pi_t + \varphi_\pi z_{\pi,t+1} + \varphi_{\pi g} \sigma_{g,t} z_{g,t+1} + \varphi_{\pi \sigma} \sqrt{q_t} z_{\sigma,t+1}. \end{aligned} \quad (11)$$

Model I includes only the autonomous inflation and constant volatility. Even with stochastic volatility, Model II still has no genuine inflation risk premium, since the inflation innovation is exogenous.¹⁵ When there is stochastic volatility, either through the growth channel (Model III) or uncertainty channel (Model IV), a genuine inflation risk premium exists and money neutrality is broken implicitly. Our preferred inflation specification (10) or Model V incorporates all three channels.¹⁶

For each of the five model specifications, one can solve for the bond yield, the bond risk premium, the predictability slope coefficient, and R^2 , when regressing the bond risk premium on the variance risk premium. We present the general result of Model V here, as others are either special cases or very easy to derive.¹⁷ The nominal bond yield can be expressed as an affine

¹⁵The ability or inability of Model II in explaining both the level of bond risk premia and the predictability of the variance risk premium is also examined by Zhou (2010).

¹⁶There is a growing literature that examines the stochastic volatility or uncertainty effect in real macroeconomic variables (see, e.g., Bloom, 2009; Bloom *et al.*, 2010; Benigno *et al.*, 2010; Fernández-Villaverde and Rubio-Ramírez, 2010).

¹⁷The analytical solutions for bond prices, bond risk premia, the predictability R^2 , slope coefficients for Models I–V, and the real economy are provided in a technical note (Zhou, 2011).

function of the state variables,

$$y_t^n = -\frac{1}{n}[A(n) \ B(n) \ C(n) \ D(n)][1 \ \sigma_{g,t}^2 \ q_t \ \pi_t]', \tag{12}$$

where the coefficients $A(n)$, $B(n)$, $C(n)$, and $D(n)$ are solutions to ordinary difference equations.

Let rx_{t+1}^{n-1} be the bond excess return from t to $t + 1$ for an n -period bond holding one period, then its expected value brp_t^n or bond risk premium is given by

$$\begin{aligned} \text{brp}_t^n &= D(n-1)\varphi_{\pi g} \left(-\frac{\theta}{\psi} + \theta - 1 - \varphi_{\pi g} \right) \sigma_{g,t}^2 \\ &\quad \times \{ [B(n-1) + D(n-1)\varphi_{\pi\sigma}] [(\theta-1)\kappa_1 A_\sigma - \varphi_{\pi\sigma}] \\ &\quad + C(n-1)(\theta-1)\kappa_1 A_q \varphi_q^2 \} q_t - D(n-1)\varphi_\pi^2. \end{aligned} \tag{13}$$

The first two items reflect consumption and uncertainty risk premia that are amplified by the endogenous inflation shock parameters $\varphi_{\pi g}$ and $\varphi_{\pi\sigma}$, while the third item captures the autonomous inflation shock rough φ_π .

Our modeling framework also has implications for the predictability pattern of the bond risk premium by the variance risk premium. In a regression $\text{brp}_t^n = a + b\text{VRP}_t$, the model-implied slope coefficient and R^2 are given by

$$\begin{aligned} b &= \frac{\text{Cov}(\text{brp}_t^n, \text{VRP}_t)}{\text{Var}(\text{VRP}_t)} = \frac{\{\cdot\}}{(\theta-1)\kappa_1 [A_\sigma + A_q \kappa_1^2 (A_\sigma^2 + A_q^2 \varphi_q^2) \varphi_q^2]}, \\ R^2 &= \frac{b^2 \text{Var}(\text{VRP}_t)}{\text{Var}(\text{brp}_t^n)} = \frac{\{\cdot\}^2 \text{Var}(q_t)}{\{\cdot\}^2 \text{Var}(q_t) + D(n-1)^2 \varphi_{\pi g}^2 \left(-\frac{\theta}{\psi} + \theta - 1 - \varphi_{\pi g} \right)^2 \text{Var}(\sigma_{g,t}^2)}, \end{aligned}$$

where $\{\cdot\} \equiv [B(n-1) + D(n-1)\varphi_{\pi\sigma}] [(\theta-1)\kappa_1 A_\sigma - \varphi_{\pi\sigma}] + C(n-1)(\theta-1)\kappa_1 A_q \varphi_q^2$. Using these two metrics, we can evaluate whether the proposed inflation dynamics can reproduce the empirical pattern of bond return predictability from the variance risk premium as presented in Sec. 2.

3.3. Calibrating bond risk premia and return predictability

The key to match the bond risk premium dynamics is through calibrating the inflation process, while leaving the choices of preference parameters and real economy dynamics similar to the existing studies (see, e.g., Bansal and Yaron, 2004; Bollerslev *et al.*, 2009). Across all five models, as seen in Panel A of Table 5, we choose the same inflation level and persistence such that the annualized inflation rate is 2.4%. The choices of the volatility parameters are such that the annualized inflation volatility is 4.5%. When there are two or

Table 5. Model calibration parameter setting.

Panel A: Inflation Dynamics					
	Model I	Model II	Model III	Model IV	Model V
Constant	$a_\pi = 8 \times 10^{-4}$	$a_\pi = 8 \times 10^{-4}$	$a_\pi = 8 \times 10^{-4}$	$a_\pi = 8 \times 10^{-4}$	$a_\pi = 8 \times 10^{-4}$
Persistence	$\rho_\pi = 0.60$	$\rho_\pi = 0.60$	$\rho_\pi = 0.60$	$\rho_\pi = 0.60$	$\rho_\pi = 0.60$
Autonomous	$\varphi_\pi = 0.0104$	$\varphi_\pi = 0.2221$	$\varphi_\pi = 0.0073$	$\varphi_\pi = 0.0073$	$\varphi_\pi = 0.006$
Consumption			$\varphi_{\pi g} = -0.1570$		$\varphi_{\pi g} = -0.1282$
Uncertainty				$\varphi_{\pi\sigma} = 0.2324$	$\varphi_{\pi\sigma} = 0.1897$

Panel B: Real Economy	
Preference	$\delta = 0.997$
	$\gamma = 2$
	$\psi = 1.5$
Endowment	$\mu_g = 0.0015$
	$a_\sigma = 0.0011$
	$\rho_\sigma = 0.5$
	$a_q = 2 \times 10^{-5}$
Uncertainty	$\rho_q = 0.98$
	$\varphi_q = 0.006$

Notes: This table reports the calibration parameter values for the real economy model, similar to Bollerslev *et al.* (2009) and Zhou (2010). The Campbell–Shiller linearization constants are $\kappa_1 = 0.9$ and hence, $\kappa_0 = 0.3251$. The inflation dynamics parameters are adapted from Gallmeyer *et al.* (2009) for our sample period of January 1990 to September 2010.

three innovations in inflation shocks, as in Models III–V, the parameters are set such that each component contributes equally to the total inflation volatility. Note that the inflation dynamics — level, persistence, and volatility — are almost the same as the exogenous process in Gallmeyer *et al.* (2009), which imposes certain disciplines on our calibration exercise. The choices of preference structure and real economy parameters, as seen in Panel B of Table 5, are largely similar to those in Zhou (2010).

Our main calibration result on short-term bond risk premia levels is reported in Panel A of Table 6. The observed bond risk premia of two- to six-month Treasury bills for a one-month holding period range from 33 to 75 basis points (bps). It is instructive to use the real bond as a benchmark — 174–338 bps, which is far exceeding the observed levels. Therefore, it does not come as a surprise that exogenous inflation, either with stochastic volatility (Model II) or without (Model I), will overshoot bond risk premia even more since the exogenous inflation shock only adds on to the bond risk premium, which is purely driven by the Jensen’s inequality term but not by a genuine risk premium effect.

Table 6. Calibrated model-implied financial market risk premia.

Panel A: Bond Risk Premia							
	Data	Real Model	Model I	Model II	Model III	Model IV	Model V
2 Month Bill	0.33	1.74	1.82	1.85	0.93	1.26	0.69
3 Month Bill	0.46	2.61	2.74	2.84	1.13	1.89	0.84
4 Month Bill	0.45	3.05	3.20	3.38	1.08	2.19	0.80
5 Month Bill	0.67	3.27	3.44	3.68	0.97	2.34	0.71
6 Month Bill	0.75	3.38	3.56	3.85	0.86	2.42	0.62

Panel B: Variance and Equity Risk Premia		
	Data	Model
Variance Premium		
Mean	18.47	10.84
Std Dev	17.48	10.34
Skewness	2.18	1.87
Kurtosis	7.78	8.04
Equity Premium		
Equity Premium	3.58	5.61
Equity Volatility	14.60	21.91
Risk-Free Rate	1.13	1.12
Risk-Free Rate Volatility	3.37	14.61

Notes: This table reports the calibration output values for bond risk premia and variance risk premia from the stochastic uncertainty model of consumption and inflation dynamics used in this paper. The observed bond risk premia and variance risk premia are from the empirical exercise of this paper for the sample period of 1990:01–2008:09.

In essence, we are facing a challenge of simultaneously matching the levels of bond risk premia and the moments of variance risk premium. As shown in Panel B of Table 6, our sample variance risk premium has a mean of 21.57 and a standard deviation of 23.42 (pre-crisis values are 18.5 and 17.5, respectively), which is within the typical range found in recent empirical studies (Bollerslev *et al.*, 2009; Drechsler and Yaron, 2011). Our real economy model can match the observed (pre-crisis) variance risk premium reasonably well with a mean of 10.84 and a standard deviation of 10.34. Our model also does a decent job matching the skewness (2.18) and kurtosis (7.78), producing model implied values of 1.87 (skewness) and 8.04 (kurtosis).

Of course, our calibration strategy for the real model, and consequently for the exogenous inflation Models I and II, is to match the variance risk premium as best as we can but sacrifice by overfitting the bond risk premia about six to seven times larger. A similar trade-off is also reported in Zhou (2010), where the real economy model or autonomous inflation Model II is able to

match well with the bond risk premia (44–86 bps in data and 73–94 bps in model), but severely undershoots the variance risk premia (18.5 in the data compared to 4.62 in model). Therefore, without dropping the money neutrality assumption implicit in the autonomous inflation dynamics, there is little hope that one can simultaneously match bond and variance risk premia.

It is interesting to note that when money neutrality is violated as in Model III, the model-implied bond risk premia can be dramatically lowered to around 86–113 bps, compared to the exogenous inflation Model II (around 185–385 bps). This improvement is primarily driven by the negative co-movement between inflation and consumption innovations ($\varphi_{\pi g} = -0.157 < 0$, Panel A of Table 5). The negative correlation between inflation and consumption shocks is consistent with more recent empirical evidence that show both growth and inflation are in a moderate range (see, e.g., Piazzesi and Schneider, 2006; Campbell *et al.*, 2009). Similarly, when the inflation shock is positively correlated with the uncertainty shock, as in Model IV ($\varphi_{\pi\sigma} = 0.1897 > 0$, Panel A of Table 5), bond risk premia also moderate to around 126–243 bps from the exogenous inflation in Model II (around 185–385 bps). Intuitively this could happen as volatility shocks — although uncorrelated with consumption shocks — are negatively correlated with market risk premia (see, e.g., Bansal and Yaron, 2004), therefore, any inflation shock which works through the uncertainty channel reduces bond risk premia through a discount rate effect.

Finally, in Model V, combining both cash flow and the uncertainty channels of inflation effects seems to produce reasonable bond risk premia — 62 to 84 bps — the closest to observed range of 33 to 75 bps. This is indeed a combined effect from lower risk premium of inflation's growth channel and lower risk premium of inflation's uncertainty channel, and as such, Model V may prove to be a more flexible way of modeling inflation risk in matching bond risk premium dynamics. Our result on the nominal risk-free rate and five-year yield is also reasonable — 4.38 and 2.93%, respectively, while the other four models produce a similar size of the risk-free rate but with a five-year yield ranging from -2.05 to negative infinity. Again this result reflects the challenge of simultaneously matching bond risk premia and variance risk premium.

Long-term bond risk premia cannot be matched well by our model that has only three underlying shocks. Model V implies bond risk premia of 3–10 bps versus the observed levels of 95–278 bps, while other models imply negative bond risk premia and some are near negative infinity. In terms of the real economy model, Panel B of Table 6, Model V produces a reasonable equity premium of 5.61% and an equity volatility of 21.91%. The model also matches the real risk-free rate — 1.12 — quite well, but the risk-free rate volatility of

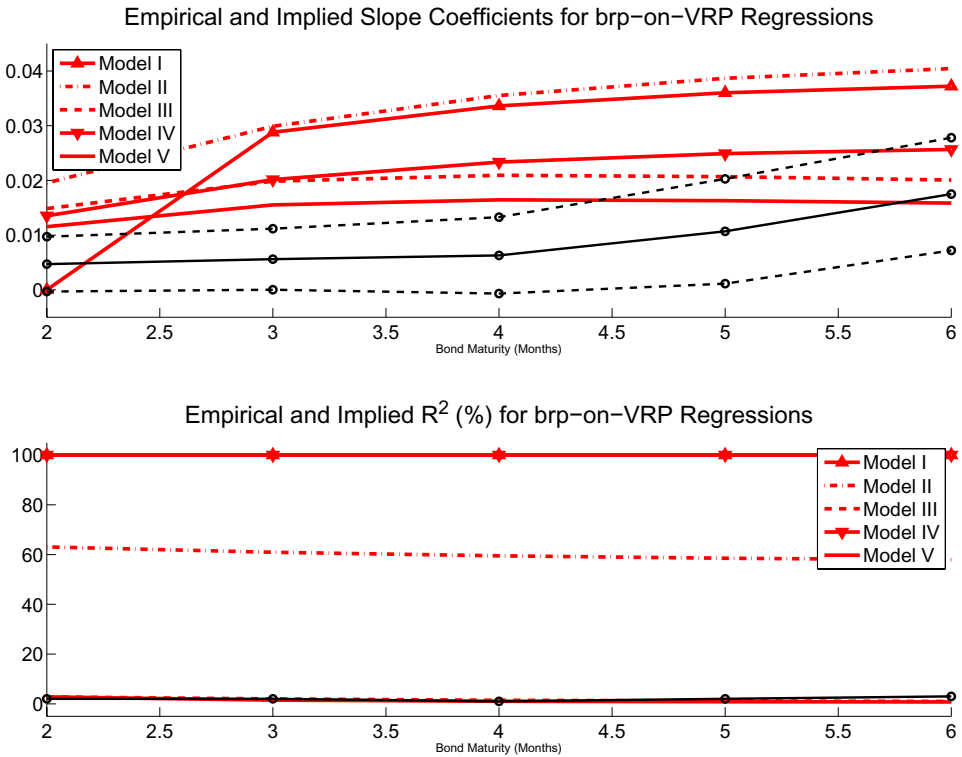


Fig. 4. Model-implied and estimated slopes and R^2 s for two- to six-month T-bills.

Notes: The figure shows the calibrated model-implied slope coefficients and R^2 s (thick lines) for regressing the two to six months Treasury bill one-month excess returns on the variance risk premium, along with their estimated empirical counterparts (thin lines with circles), and 95% confidence bands of the slope coefficients (thin dashed lines with circles).

14.61% is much higher than the historical average of around 3.37%. The overshooting of the risk-free rate volatility and underfitting of long-term bond risk premia are the closely related outcomes of limiting the setup to only three risk factors.

The model-implied predictability regression slope coefficients and R^2 s are plotted in Fig. 4, along with the empirical estimated ones. As shown in the top panel, the predictability slope coefficients of one-month excess bond return regressions clearly show a gradual upward trend for 2–6 month Treasury bills. Model I and Model II overfit the predictability slopes by quite a large margin. Models III and IV improve significantly and fall within the 95% confidence bands for 5–6 month t-bills. Our preferred Model V seems to fit reasonably well with the slope coefficients and is the closest to the 95% confidence bands, in fact it matches the 5–6 month t-bills almost exactly. In

the bottom panel, Models I and VI produce predictability R^2 s of 100% by construction, since both bond risk premium and variance risk premium are driven by the same uncertainty factor, q_t , alone (Zhou, 2010). Model II seems to improve significantly as the bond risk premium also loads on the consumption growth risk $\sigma_{g,t}^2$. The R^2 's implied by Models III and V actually get very close to the observed ones.

In summary, our preferred Model V, which incorporates inflation's exposure to both growth and uncertainty risks, seems to have the potential to match both the observed bond risk premia levels and the predictability pattern from the variance risk premium.

4. Conclusion

This paper documents the predictive power of the equity market variance risk premium for bond excess returns. The variance risk premium is defined as the difference between the risk-neutral and objective expectations of return variations and it is estimated without a forward-looking bias. The predictive power is shown to be particularly strong in the short-run for a one-month horizon, irrespective of the underlying bond maturities. The information contained in the variance risk premium is orthogonal to other known predictors of bond risk premia — forward rates, jump risk, and macro variables. The previously documented predictors for bond risk premia are particularly powerful for longer horizons. Short-term bond risk premia exhibit pronounced spikes around major economic and financial crises, which is in contrast to the cyclical swings typically observed in long-term bond risk premia.

We then propose a model that features time-varying uncertainty about real and nominal quantities, along with investors' preferences for early resolution for uncertainty, to produce the level and predictability of bond risk premia, which have been found difficult to pin down in standard asset pricing models. While the real side of the economy follows earlier literature, the inflation process consists of two key ingredients — one stochastic volatility process that covaries with the consumption growth and the other that covaries with the consumption uncertainty, which gives rise to a genuine inflation risk premium.

In our calibration exercise, the model implied bond risk premia are only several basis points away from their empirical counterparts. The model also produces a reasonable equity premium, risk-free rate, and equity volatility but overshoots the risk-free rate volatility. The average variance risk

premium produced by our model is only slightly lower than the estimated empirical variance risk premium. In addition, the higher order moments (variance, skewness and kurtosis) of the variance risk premium are also fitted quite nicely. Finally, the model is able to replicate the predictive power of the equity variance risk premium for bond risk premia remarkably well.

Acknowledgments

We would like to thank Tim Bollerslev, Mike Chernov, Frank Diebold, Urban Jermann, George Tauchen, and seminar participants at Lund, Duke, LSE, and HKUST for their helpful comments. Mueller and Vedolin gratefully acknowledge the financial support from STICERD and the Systemic Risk Centre at LSE, and the British Academy.

References

- Baele, L., G. Bekaert, and K. Inghelbrecht, 2010, The Determinants of Stock and Bond Return Comovements, *Review of Financial Studies* 23, 2374–2428.
- Bansal, R., and I. Shaliastovich, 2009, Learning and asset-price jumps, Working Paper.
- Bansal, R., and I. Shaliastovich, 2010, A Long-Run Risks Explanation of Predictability Puzzles in Bond and Currency Markets, Working Paper.
- Bansal, R., and A. Yaron, 2004, Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles, *Journal of Finance* 59, 1481–1509.
- Bekaert, G., and E. Engstrom, 2009, Asset Return Dynamics Under Bad Environment-Good Environment Fundamentals, Working Paper.
- Bekaert, G., M. Hoerova, and M. Lo Duca, 2010, Risk, Uncertainty and Monetary Policy, Working Paper NBER 16397.
- Benigno, P., L. A. Ricci, and P. Surico, 2010, Unemployment and Productivity in the Long Run: The Role of Macroeconomic Volatility, Working Paper.
- Bloom, N., M. Floetotto, and N. Jaimovich, 2010, Really Uncertain Business Cycles, Working Paper.
- Bloom, N., 2009, The Impact of Uncertainty Shocks, *Econometrica* 77, 623–685.
- Bollerslev, T., G. Tauchen, and H. Zhou, 2009, Expected Stock Returns and Variance Risk Premia. *Review of Financial Studies* 22, 4463–4492.
- Bollerslev, T., and V. Todorov, 2011, Tails, Fears, and Risk Premia, *Journal of Finance* 66, 2165–2211.
- Bollerslev, T., and H. Zhou, 2007, Expected Stock Returns and Variance Risk Premia, Finance and Economics Discussion Series 2007–11, Federal Reserve Board.
- Britten-Jones, M., and A. Neuberger, 2000, Option Prices, Implied Price Processes, and Stochastic Volatility, *Journal of Finance* 55, 839–866.
- Buraschi, A., and P. Whelan, 2010, Macroeconomic Uncertainty, Difference in Beliefs, and Bond Risk Premia, Working Paper.

- Campbell, J. Y., and R. J. Shiller, 1991, Yield Spreads and Interest Rate Movements: A Bird'S Eye View, *Review of Economic Studies* 58, 495–514.
- Campbell, J. Y., and R. J. Shiller, 1988, The Dividend-price Ratio and Expectations of Future Dividends and Discount Factors, *Review of Financial Studies* 1, 195–228.
- Campbell, J. Y., A. Sunderam, and L. M. Viceira, 2009, Inflation Bets or Deflation Hedges? The Changing Risks of Nominal Bonds, Working Paper, Harvard University.
- Carr, P., and L. Wu, 2009, Variance Risk Premiums, *Review of Financial Studies* 22, 1311–1341.
- CBOE, 2003, Vix: Cboe Volatility Index, CBOE White Paper.
- Cochrane, J. H., and M. Piazzesi, 2005, Bond Risk Premia, *American Economic Review* 95, 138–160.
- Corsi, F., 2009, A Simple Approximate Long-Memory Model of Realized Volatility, *Journal of Financial Econometrics* 7, 174–196.
- Corsi, F., D. Pirino, and R. Renò, 2010, Threshold Bipower Variation and The Impact of Jumps on Volatility Forecasting, *Journal of Econometrics* 156, 276–288.
- Demeterfi, K., E. Derman, M. Kamal, and J. Zou, 1999, A Guide to Volatility and Variance Swaps, *Journal of Derivatives* 6, 9–32.
- Doh, T., 2010, Long Run Risks in the Term Structure of Interest Rates: Estimation, Working Paper, Kansas City Fed RWP 08-11.
- Drechsler, I., and A. Yaron, 2011, What's Vol got to do With It, *Review of Financial Studies* 24, 1–45.
- Duffee, G., 2011, Information in (and not in) the Term Structure, *Review of Financial Studies* 24, 2895–2934.
- Duffee, G. R., 1996, Idiosyncratic Variation of Treasury Bill Yields, *Journal of Finance* 51, 527–551.
- Fama, E. F., and R. R. Bliss, 1987, The Information in Long-Maturity Forward Rates, *American Economic Review* 77, 680–692.
- Fernández-Villaverde, J., and J. Rubio-Ramírez, 2010, Macroeconomics and volatility: Data, Models, and Estimation, NBER Working Paper 16618.
- Fontaine, J.-S., and R. Garcia, 2010, Bond liquidity premia, Working Paper.
- Gabaix, X., 2009, Variable Rare Disasters: An Exactly Solved Framework for Ten Puzzles In Macro-Finance, Working Paper, NYU Stern.
- Gallmeyer, M., B. Hollifield, F. Palomino, and S. Zin, 2009, Term premium dynamics and the Taylor Rule, Working Paper.
- Greenwood, R., and D. Vayanos, 2010, Bond Supply and Excess Bond Returns, Working Paper.
- Gürkaynak, R., B. Sack, and J. Wright, 2007, The U.S. Treasury Yield Curve: 1961 to the Present, *Journal of Monetary Economics* 54, 2291–2304.
- Hasseltoft, H., 2010, Stocks, Bonds, and Long-run Consumption Risks, Working Paper.
- Huang, J., and Z. Shi, 2010, Determinants of Bond Risk Premia, Working Paper.
- Jiang, G. J., and Y. S. Tian, 2005, Model-free Implied Volatility and its Information Content, *Review of Financial Studies* 18, 1305–1342.
- Koijen, R. S. J., H. N. Lustig, and S. Van Nieuwerburgh, 2010, The Cross-Section and Time-Series of Stock and Bond Returns, Working Paper, University of Chicago.

P. Mueller, A. Vedolin & H. Zhou

- Ludvigson, S. C., and S. Ng, 2009, Macro Factors in Bond Risk Premia, *Review of Financial Studies* 22, 5027–5067.
- Ludvigson, S. C., and S. Ng, 2010, A Factor Analysis of Bond Risk Premia, in A. Ulah and D. E. A. Giles (editors), *Handbook of Empirical Economics and Finance*, Vol. 1, Chapman and Hall.
- Patton, A., and K. Sheppard, 2011, Good Volatility, Bad Volatility: Signed Jumps and the Persistence of Volatility, Working Paper.
- Pennacchi, G., 1991, Identifying the Dynamics of Real Interest Rates and Inflation: Evidence Using Survey Data, *Review of Financial Studies* 4, 53–86.
- Piazzesi, M., and M. Schneider, 2006, Equilibrium Yield Curve, *NBER Macroeconomics Annual* 21, 389–472.
- Stock, J. H., and M. W. Watson, 2002, Macroeconomic Forecasting Using Diffusion Indexes, *Journal of Business and Economic Statistics* 20, 147–162.
- Sun, T.-S., 1992, Real and Nominal Interest Rate: A Discrete-Time Model and its Continuous-Time Limit, *Review of Financial Studies* 5, 581–611.
- Vayanos, D., and J.-L. Vila, 2009, A Preferred-Habitat Model of the Term-Structure of Interest Rates, Working Paper.
- Wachter, J. A., 2006, A Consumption-Based Model of the Term Structure of Interest Rates, *Journal of Financial Economics* 79, 365–399.
- Whaley, R. E., 2000, The Investor Fear Gauge, *Journal of Portfolio Management* 26, 12–17.
- Wright, J. H., and H. Zhou, 2009, Bond Risk Premia and Realized Jump Risk, *Journal of Banking and Finance* 33, 2333–2345.
- Wu, S., 2008, Long-Run Consumption Risk and the Real Yield Curve, Working Paper, Department of Economics, University of Kansas.
- Xiong, W., and H. Yan, 2010, Heterogeneous Expectations and Bond Markets, *Review of Financial Studies* 23, 1433–1466.
- Zhou, H., 2010, Variance Risk Premia, Asset Predictability Puzzles, and Macroeconomic Uncertainty, Working Paper, Federal Reserve Board.
- Zhou, H., 2011, Term Structure of Interest Rates with Inflation Uncertainty, Technical Report, Federal Reserve Board.