

Stock Return Predictability and Variance Risk Premia: Statistical Inference and International Evidence

Tim Bollerslev, James Marrone, Lai Xu, and Hao Zhou*

Abstract

Recent empirical evidence suggests that the variance risk premium predicts aggregate stock market returns. We demonstrate that statistical finite sample biases cannot “explain” this apparent predictability. Further corroborating the existing evidence of the United States, we show that country-specific regressions for France, Germany, Japan, Switzerland, the Netherlands, Belgium, and the United Kingdom result in quite similar patterns. Defining a “global” variance risk premium, we uncover even stronger predictability and almost identical cross-country patterns through the use of panel regressions.

I. Introduction

A number of recent studies have argued that aggregate U.S. stock market return is predictable over horizons ranging up to two quarters based on the difference between option-implied and actual realized variation measures, or the so-called variance risk premium (see Bollerslev, Tauchen, and Zhou (BTZ) (2009), Drechsler and Yaron (2011), Gabaix (2012), Kelly (2011), Zhou (2010), and Zhou

*Bollerslev, bolle@duke.edu, Department of Economics, Duke University, Box 90097, Durham, NC 27708, and NBER and CREATES; Marrone, jmar@uchicago.edu, Department of Economics, University of Chicago, 1126 E 59th St, Chicago, IL 60637; Xu, lxu100@syr.edu, Whitman School of Management, Syracuse University, 721 University Ave, Syracuse, NY 13244; and Zhou, zhouh@pbcfs.tsinghua.edu.cn, PBC School of Finance, Tsinghua University, 43 Chengfu Rd, Haidian District, Beijing, 100083, P. R. China. We thank Geert Bekaert, Stephen Brown (the editor), Wayne Ferson, Bryan Kelly (the referee), Qianqiu Liu, Anthony Neuberger, Andrew Patton, Christian Schlag, George Tauchen, Liuren Wu, Guofu Zhou, and conference and seminar participants at the 2011 European Summer Symposium in Financial Markets (Asset Pricing) in Gerzensee, 2011 New York University Stern Volatility Institute Conference, 2011 Hedge Fund Conference at Imperial College London, 2011 China International Conference in Finance (CICF) in Wuhan, 2011 National Bureau of Economic Research (NBER)-National Science Foundation (NSF) Time Series Conference at Michigan State University, 2011 Inquire Europe Autumn Seminar in Luxembourg, Notre Dame University, University of Zurich, and the 2011 Duke Financial Econometrics Lunch Group for their helpful comments. We also gratefully acknowledge the 2011 CICF Best Paper Award. Bollerslev’s research was supported by a grant from the NSF to the NBER and Center for Research in Econometric Analysis of Time Series (CREATES) funded by the Danish National Research Foundation. This paper was written while Zhou was at the Board of Governors of the Federal Reserve System. The analysis and conclusions set forth are those of the authors and do not indicate concurrence by other members of the research staff or the Board of Governors.

and Zhu (2009), among others). These findings are distinctly different from the longer-run multiyear return predictability patterns that have been studied extensively in the existing literature, in which the predictability is typically associated with more traditional valuation measures such as dividend yields, price-to-earnings (PE) ratios, or consumption-wealth ratios (see Fama and French (1988), Campbell and Shiller (1988), and Lettau and Ludvigson (2001), among others). The present paper builds and further expands on the scope of these striking new empirical findings.

The variance risk premium is formally defined as the difference between the risk-neutral and statistical expectations of the future return variation.¹ It may be interpreted as a measure of both aggregate risk aversion and aggregate economic uncertainty. In our main empirical investigations reported on below, we follow BTZ (2009) in approximating the variance risk premium by the difference between 1-month forward-looking model-free option-implied variances and the actual 1-month realized variances at the time. This directly observable proxy has the obvious advantage of being simple to implement and completely model free.

Our investigations are essentially threefold. First, to assess the validity of the previously documented predictability patterns, we report the results from a Monte Carlo simulation study designed to closely mimic the dynamic dependencies inherent in daily U.S. returns and variance risk premia. Our results clearly show that statistical biases can not “explain” the documented return predictability patterns.

Second, in a separate effort to corroborate and further expand on the existing empirical evidence based on monthly U.S. data prior to the advent of the financial crisis, we extend the same basic return predictability regressions to seven other countries and more recent “out-of-sample” data spanning the financial crisis. We show that the same predictability pattern that exists for the United States holds true for most of the other countries, although the magnitude in each is somewhat attenuated.

Third, motivated by this apparent commonality across countries, we define a “global” variance risk premium. We show that this simple aggregate worldwide variance risk premium results in strong predictability for all of the countries in the sample.

The finite sample properties of overlapping long-horizon return regressions have been studied extensively in the literature. Boudoukh, Richardson, and Whitelaw (2008), for instance, have recently shown that even in the absence of any increase in true predictability, the values of the R^2 s in regressions involving highly persistent predictor variables and overlapping returns, by construction, will increase roughly proportionally to the return horizon and the length of the overlap.² By contrast, the variance risk premium is not especially persistent at the monthly horizon. Our simulations are based on a bivariate vector autoregressive (VAR)-generalized autoregressive conditional heteroskedasticity

¹The variance risk premium is sometimes defined the other way around as the statistical minus risk-neutral expectations. This, of course, is immaterial for all of the results reported on below.

²Closely related issues pertaining to the use of persistent predictor variables have also been studied (see Stambaugh (1999), Ferson, Sarkissian, and Simin (2003), Baker, Taliaferro, and Wurgler (2006), Campbell and Yogo (2006), Ang and Bekaert (2007), and Goyal and Welch (2008), among others).

(GARCH)-dynamic conditional correlation (DCC) model designed to closely mimic the relevant joint dynamic dependencies in the daily return and the variance risk premium. We find that the robust t -statistics usually employed in the literature are reasonably well behaved, albeit slightly oversized under the null hypothesis of no predictability. We also find that the quantiles in the finite sample distribution of the R^2 s from the regressions spuriously increase with the return horizon under the null of no predictability, and are distinctly different from the hump-shaped R^2 s actually observed in the U.S. data at the 1- to 12-month horizons.

Guided by the Monte Carlo simulations, we rely on simple ordinary least squares (OLS) regressions along with Newey-West (NW) (1987) based t -statistics based on simulated critical value to summarize our new international evidence. Due to data availability and liquidity considerations, we restrict our attention to the eight financial markets of France, Germany, Japan, Switzerland, the Netherlands, Belgium, the United Kingdom, and the United States. Regressing the individual country returns on the country-specific variance risk premia results in similar hump-shaped regression coefficients and R^2 s for all eight countries. However, the degree of predictability afforded by the country-specific variance risk premia and the statistical significance of the regression coefficients are generally not as strong as the previously reported results for the United States.

These results naturally raise the question of whether worldwide variance risk, as opposed to country-specific variance risk, is being priced by the market? To investigate this idea, we construct a simple global variance risk premium proxy, defined as the market capitalization weighted average of the individual country variance risk premia. Restricting the effect on this global variance risk premium to be the same across countries in a panel return regression results in much stronger findings for all of the countries, with a systematic peak in the degree of predictability around the 4-month horizon. Moreover, the degree of predictability afforded by this global variance risk premium easily exceeds that of the implied and realized variation measures when included in isolation. It also clearly dominates that of other traditional predictor variables that have been shown to work well over longer annual horizons, including the PE ratio.³

Our use of the variance difference as a simple proxy for the variance risk premium implicitly assumes that the volatility follows a random walk.⁴ To investigate the sensitivity of our main international findings to this simplifying assumption, we define a forward-looking global variance risk premium from the differences between the individual countries' 1-month option-implied variance and the corresponding 1-month VAR-based forecasts for the actual variance. This alternative definition of the global variance risk premium gives rise to almost identical international return predictability patterns.

³Related evidence has also been reported in a few other recent studies pertaining to other markets. In particular, in concurrent independent work, Londono (2011) finds that the U.S. variance risk premium predicts several foreign stock market returns. In a slightly different context, Mueller, Vedolin, and Zhou (2011) argue that the U.S. variance risk premium predicts bond risk premia, beyond the predictability afforded by forward rates, while Buraschi, Trojani, and Vedolin (2014) and Zhou (2010) show that the variance risk premium also helps predict credit spreads, over and above the typical interest rate predictor variables.

⁴Of course, the variance difference may simply be interpreted as a powerful predictor variable in its own right.

Putting things into perspective, our new empirical findings are clearly related to the large existing literature on international stock return predictability (see Harvey (1991), Bekaert and Hodrick (1992), Campbell and Hamao (1992), and Ferson and Harvey (1993), among others). However, the focus of this literature has traditionally been on longer-run multiyear return predictability. By contrast, our results pertaining to the global variance risk premium concern much shorter runs within year predictability and are essentially “orthogonal” to the findings reported in the existing literature.⁵

The remainder of the paper is organized as follows: Section II presents our simulation-based results pertaining to the statistical inference procedures and robustness of the existing empirical evidence for the United States. Section III details our international data and country-specific return regressions. Section IV discusses the results based on our new global variance risk premium and the combined panel regressions for all of the countries. Section V provides conclusions.

II. General Setup and Monte Carlo Simulations

The key empirical findings reported in BTZ (2009), and the subsequent studies cited above, are based on simple OLS regressions of the returns on the aggregate market portfolio over monthly and longer return horizons on a measure of the 1-month variance risk premium.

In particular, let $r_{t,t+\tau}$ and VRP_t denote the continuously compounded return from time t to time $t + \tau$ and the variance risk premium at time t , respectively. Defining the unit time interval to be 1 month, the multiperiod return regressions in BTZ (2009) may then be expressed as

$$(1) \quad \frac{1}{h} \sum_{j=1}^h r_{t+(j-1),t+j} = a(h) + b(h)VRP_t + u_{t,t+h},$$

where $t = 1, 2, \dots, T - h$ refers to the specific observations used in the regression.

Meanwhile, it is well known that in the context of overlapping return observations, the regression in equation (1) can result in spuriously large and highly misleading regression R^2 s, say $R^2(h)$, as the horizon h increases (see, e.g., the discussion and many references in Campbell, Lo, and MacKinlay (1997)). Similarly, the standard errors for the OLS estimates designed to take account of the serial correlation in $u_{t+h,t}$ based on the Bartlett kernel advocated by NW (1987), and the modification proposed by Hodrick (HD) (1992), can also both result in t -statistics for testing hypotheses about $a(h)$ and $b(h)$ that are poorly approximated by a standard normal distribution.

Most of the existing analyses pertaining to these and other related finite sample biases, however, have been calibrated to situations with a highly persistent predictor variable, as traditionally used in long-horizon return regressions. Even though the variance risk premium is fairly persistent at the daily frequency,

⁵Other recent studies highlighting short-run international predictability include Rapach, Strauss, and Zhou (2013) based on lagged U.S. returns, Ang and Bekaert (2007) and Hjalmarsson (2010) based on short-term interest rates, and Bakshi, Panayotov, and Skoulakis (2011) based on the Baltic Dry Index.

it is much less so at the monthly level, and as such, one might naturally expect the finite sample biases to be less severe in this situation.⁶ Our Monte Carlo simulations discussed in the next section confirm this conjecture in an empirically realistic setting designed to closely mimic the joint dependencies in actual daily returns and variance risk premia.

A. Simulation Design

The model underlying our simulations is based on daily Standard & Poor’s (S&P) 500 composite index returns (obtained from the Center for Research in Security Prices (CRSP)). The corresponding daily observations on the variance risk premium are defined as $VRP_t = IV_t - RV_{t-1,t}$, where we rely on the square of the new VIX index (obtained from the Chicago Board Options Exchange (CBOE)) to quantify the implied variation IV_t , and the summation of current and previous 20 trading days’ daily realized variances (obtained from S&P) together with the squared overnight returns to quantify the total realized variation over the previous month $RV_{t-1,t}$.⁷

The sample period runs from Feb. 1, 1996, to Dec. 31, 2007, for a total of 2,954 daily observations. The end of the sample purposely coincides with that in BTZ (2009). Later, we will investigate the sensitivity of the empirical results to the inclusion of more recent data involving the financial crisis. The span of the data exactly matches the length of the commonly available sample for the eight countries we analyze.

After some experimentation, we arrived at the following bivariate VAR(1)-GARCH(1, 1)-DCC model (see Engle (2002) for additional details on the DCC model) for the two daily time series, corresponding to $\Delta = 1/20$,

$$\begin{aligned}
 r_{t-\Delta,t} &= -1.958e-5 - 0.009r_{t-2\Delta,t-\Delta} + 0.025VRP_{t-\Delta} + \epsilon_{t,r}, \\
 &\quad (0.001) \quad (0.016) \quad (0.010) \\
 VRP_t &= 3.759e-5 + 0.033r_{t-2\Delta,t-\Delta} + 0.972VRP_{t-\Delta} + \epsilon_{t,VRP}, \\
 &\quad (0.001) \quad (0.017) \quad (0.010) \\
 \sigma_{t,r}^2 &= 1.280e-6 + 0.071\epsilon_{t-\Delta,r}^2 + 0.920\sigma_{t-\Delta,r}^2, \\
 &\quad (1.68e-6) \quad (0.004) \quad (0.008) \\
 \sigma_{t,VRP}^2 &= 2.038e-7 + 0.133\epsilon_{t-\Delta,VRP}^2 + 0.871\sigma_{t-\Delta,VRP}^2, \\
 &\quad (7.59e-6) \quad (0.004) \quad (0.028) \\
 Q_t &= \begin{pmatrix} 0.997 & -0.754 \\ (0.036) & (0.040) \\ -0.754 & 1.023 \\ (0.040) & (0.060) \end{pmatrix} + 0.011\eta_{t-\Delta}\eta'_{t-\Delta} + 0.979Q_{t-\Delta}, \\
 R_t &= \text{diag}\{Q_t\}^{-1}Q_t\text{diag}\{Q_t\}^{-1},
 \end{aligned}$$

⁶The first-order autocorrelation coefficient for the monthly U.S. variance risk premium analyzed in the empirical section below equals 0.39, and it is even lower for all of the other countries included in our subsequent analysis. By comparison, the first-order autocorrelations for monthly dividend yields, PE ratios, and other valuation ratios typically employed in the long-horizon regression literature are around 0.95–0.99.

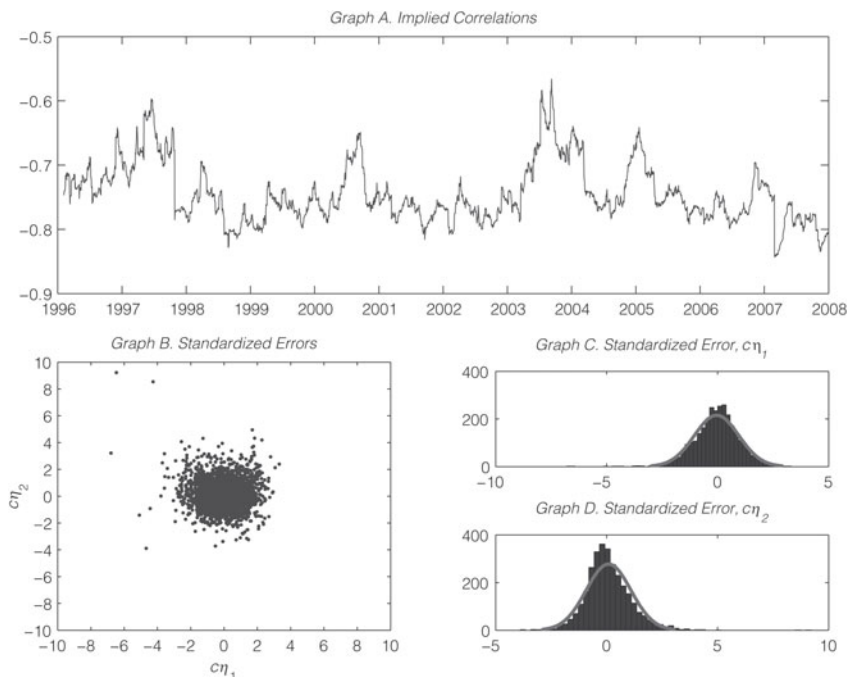
⁷This directly mirrors the definition of the variance risk premium employed in BTZ (2009). Forward-looking measures of VRP_t that align IV_t with a measure of the expected volatility $E_t(RV_{t,t+1})$ have also been used in the literature. However, this requires additional modeling assumptions for calculating $E_t(RV_{t,t+1})$, whereas the VRP_t used here has the obvious advantage of being directly observable at time t . We will return to this issue in Section IV.

where $\eta_t \equiv (\epsilon_{t,r}/\sigma_{t,r}, \epsilon_{t,VRP}/\sigma_{t,VRP})'$, and $E_{t-\Delta}(\eta_t) = 0$ and $E_{t-\Delta}(\eta_t\eta_t') = R_t$ by assumption. The specific parameter values refer to quasi-maximum likelihood estimates (QMLE) obtained under the auxiliary assumption of conditional normality, with robust standard errors following Bollerslev and Wooldridge (1992) in parentheses. With the exception of the lagged daily returns, most of the dynamic coefficients are highly significant at conventional levels.

The model implies a strong negative (on average) correlation between the innovations to the return and VRP equations. This, of course, is consistent with the well-documented “leverage” effect (see, e.g., Bollerslev, Sizova, and Tauchen (2012) and the many references therein). At the same time, as is evident from the equation for Q_t , the conditional correlation clearly varies over time and, as shown in Graph A of Figure 1, reaches a low of close to -0.85 toward the end of the sample. Graphs B–D indicate that the distribution of the estimated standardized

FIGURE 1
Estimated VAR-GARCH-DCC Model

Graph A of Figure 1 plots the daily conditional correlations between the returns and the variance risk premium implied by the estimated VAR(1)-GARCH(1,1)-DCC model described in the main text. Graphs B, C, and D provide a scatterplot and histograms, respectively, for the standardized residuals from the estimated model, $\widehat{c\eta}_t$. The daily sample used in estimating the model spans the period from Feb. 1, 1996, to Dec. 31, 2007, for a total of 2,954 daily observations.



residuals from the model (i.e., $\widehat{c\eta}_t \equiv \widehat{F}_t^{-1}\widehat{\eta}_t$, where $\widehat{F}_t \times \widehat{F}_t' = \widehat{R}_t$) are well behaved and centered at 0, with variances close to unity, albeit not normally distributed.⁸

⁸The sample means for $\widehat{c\eta}_{t,1}$ and $\widehat{c\eta}_{t,2}$ equal -0.044 and 0.088 , the standard deviations equal 0.999 and 1.007 , and the skewness and kurtosis equal -0.469 and 0.894 , and 4.913 and

Thus, all in all, the model provides a reasonably good fit to the joint dynamic dependencies inherent in the two daily series.

As such, we will use this relatively simple-to-implement model as our basic data-generating process for the Monte Carlo simulations, our analysis of the finite sample properties of the NW (1987) and HD (1992) t -statistics, and $R^2(h)$ s from the overlapping return regressions in equation (1).⁹ Our simulated finite sample distributions will be based on a total of 2,000 bootstrapped replications from the model. We will look at monthly sample frequencies and return horizons h ranging up to 12 months. The number of observations for each of the simulated samples is fixed at 149 months, corresponding to the length of the actual sample used in the estimation of the VAR-GARCH-DCC model above.¹⁰ We begin with a discussion of the size and power properties of the two t -statistics.

B. Size and Power

Our characterization of the distributions under the null hypothesis of no return predictability is based on restricting the coefficients associated with $r_{t-2\Delta, t-\Delta}$ and $VRP_{t-\Delta}$ in the return equation to be identically equal to 0, leaving all of the other coefficients at their estimated values. Panel A of Table 1 reports the resulting simulated 95th percentiles of the t^{NW} and t^{HD} test statistics, along with the regression R^2 s. In line with the evidence in the existing literature, both of

TABLE 1
Simulated Size, Power, and R^2

	Horizon							
	1	2	3	4	5	6	9	12
<i>Panel A. Simulated Size and R^2</i>								
t^{NW}	2.2602	2.5199	2.7876	2.9413	3.2413	3.2200	3.3143	3.5087
t^{HD}	2.2763	2.1871	2.0835	2.1063	2.1024	2.1237	2.1631	2.1857
Adj. R^2	3.0169	4.8366	5.7740	6.3148	7.4592	7.5017	8.1923	8.6792
<i>Panel B. Simulated Power</i>								
PW ^{NW}	0.8865	0.8450	0.7680	0.6855	0.5625	0.5070	0.3680	0.2770
PW ^{HD}	0.8070	0.7625	0.7105	0.6265	0.5470	0.4970	0.3500	0.3025

7.860, respectively. Further diagnostic checks also reveal that while the residuals from the return equation appear close to serially uncorrelated, there is some evidence for neglected longer-run serial dependencies in the equation for the variance risk premium.

⁹The bandwidth in the Bartlett kernel employed in our implementation of the NW (1987) standard errors is set to $m = \lfloor h + 4 \times ((T - hs) / 100)^{2/5} \rfloor$, where $\lfloor \cdot \rfloor$ refers to the integer value. We also experimented with the reverse regression technique suggested by HD (1992) for testing $b_s(h) = 0$. The results, available from the authors, were very similar to those for the HD t -statistic reported below.

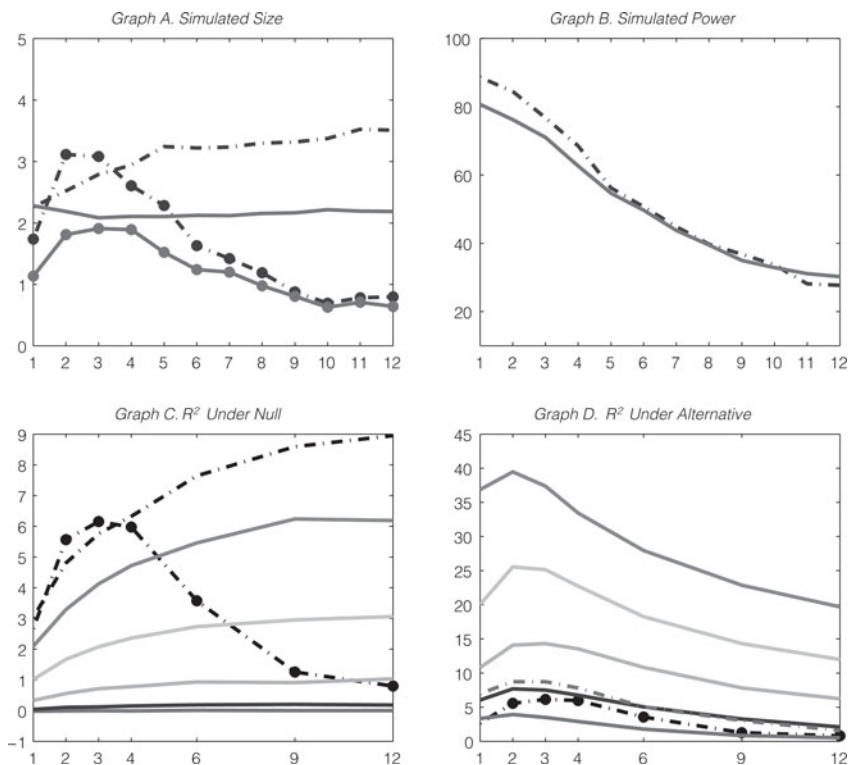
¹⁰As previously noted, this also mimics the length of the commonly available sample for the international data analyzed below.

the t -statistics exhibit nontrivial size distortions relative to the nominal one-side 95% critical value of 1.645. Also, the distortions tend to increase with the return horizon h . Moreover, consistent with the results reported in HD (1992), the biases for the NW (1987) based standard error calculations generally exceed those for the HD standard errors, and markedly more so the longer the return horizon.

To illustrate the results, we plot in Graph A of Figure 2 the simulated 95% critical values for the t^{NW} (dashed lines) and the t^{HD} (solid lines) statistics for monthly sampled data. We also include in the figure the t -statistics obtained by running these same regressions on the monthly data over the Feb. 1996 through Dec. 2007 sample period used in calibrating the simulated model. As the figure shows, the actual t^{NW} -statistics exceed the simulated critical values for return horizons in the range of 2–3 months. Meanwhile, the t^{HD} -statistics generally do not exceed the simulated critical values and, accordingly, do not support the idea of return predictability.

FIGURE 2
Simulated Size and Power

Graph A of Figure 2 reports the 95th percentiles in the finite-sample distributions of the t^{NW} (dashed line) and t^{HD} (solid line) based on simulated "monthly" data from the restricted VAR-GARCH-DCC model under the null of no predictability. The dashed-dotted lines refer to the corresponding t -statistics for actual monthly U.S. S&P 500 returns spanning Feb. 1996 to Dec. 2007. Graph C plots the quantiles in the finite-sample distribution of the R^2 from the return regression in equation (1) and simulated monthly data from the restricted VAR-GARCH-DCC model under the null of no predictability. For Graphs C and D, the dashed-dotted line refers to the corresponding R^2 's in actual daily U.S. S&P 500 returns spanning Feb. 1, 1996, to Dec. 31, 2007. Graphs B and D are based on the unrestricted VAR-GARCH-DCC model: Graph B gives the simulated monthly percentage power and the size-adjusted 5% t^{NW} (dashed line) and t^{HD} (solid line) statistics; Graph D reports the quantiles in the simulated finite-sample distribution.



In order to better understand this discrepancy in the conclusions drawn from the two tests, we report in Panel B of Table 1 the power of the tests to detect predictability implied by the unrestricted VAR-GARCH-DCC model. To facilitate comparisons, we report the size-adjusted power only for a 5% test. Not surprisingly, the power of both tests decreases with the return horizon. However, the power of the t^{NW} test exceeds that of the t^{HD} test for return horizons less than a year, and the differences appear most pronounced at the 1- to 4-month horizons. These differences are also evident in Graph B of Figure 2, which shows the plots of the relevant power curves.

In addition to the t -statistics associated with the $b(h)$ coefficients, the $R^2(h)$ s from the return regressions are also commonly used to assess the strength of the relationship and the effectiveness of the predictor variable across different horizons. Of course, it is well known that the biases exhibited by the t -statistics in the context of long-horizon return regressions with persistent predictor variable carry over to the $R^2(h)$ s, and that these need to be carefully interpreted as well (see, e.g., the aforementioned study by Boudoukh et al. (2008) for a recent analysis, along with the many references therein).

The corresponding columns in Table 1 show that, while less dramatic than the biases that exist over multiyear return horizons with highly persistent predictor variables, the $R^2(h)$ s can still be quite different from 0 under the null of no predictability in the present setting. In particular, the 95th percentiles are around 5%–6% at the 2- to 4-month horizon.

Furthermore, to this effect, we show in Graph C of Figure 2 select quantiles in the simulated distribution of the $R^2(h)$ s from daily regression that are obtained in the absence of any predictability. Consistent with the findings in the extant literature pertaining to monthly observations and longer return horizons, all of the quantiles increase monotonically with the return horizon, and this increase is especially marked for the higher percentiles. Intuitively, as the horizon increases, the overlapping return regressions become closer to a spurious-type regression.

In addition to the simulated quantiles, Graph D of Figure 2 also shows the $R^2(h)$ s obtained from the monthly return regressions implied by the same VAR-GARCH-DCC model. Comparing the actual $R^2(h)$ s to the simulated percentiles again suggests that the degree of predictability is most significant at the intermediate 2- to 4-month horizon. This, of course, is directly in line with the inference based on the t -statistics discussed in the previous section. It also supports the prior empirical evidence reported in BTZ (2009).

The hump-shaped pattern in the actual $R^2(h)$ s, with an apparent peak in the degree of predictability at the intermediate 2- to 4-month horizon, also closely mimics the patterns in the simulated quantiles for the estimated VAR-GARCH-DCC model depicted in Graphs C and D of Figure 2. Interestingly, this striking similarity arises in spite of the fact that the simulated model involves only first-order dynamics in the equations that describe the daily conditional means.

Taken as a whole, our Monte Carlo simulations and the new regression results based on daily U.S. returns discussed above clearly support the variance risk premium as a powerful predictor at the 2- to 4-month horizons. At the same time, the overlapping nature of the return regressions tends to attenuate the strength of the predictability somewhat. Hence, in an effort to further corroborate the existing

empirical evidence pertaining exclusively to the U.S. market and data prior to the 2008 financial crisis, we next turn to a discussion of our new empirical findings involving more recent data and several other countries. For each country considered, we will base our empirical investigations on monthly predictive regression and NW (1987) based standard errors with simulated critical values.

III. International Evidence

Motivated by the Monte Carlo simulation results, we will rely on the common benchmark monthly OLS regressions, along with the simulated NW (1987) critical values and t^{NW} -statistics for characterizing the return predictability internationally, keeping in mind the finite sample biases documented in the simulations. We will restrict our analysis to France, Germany, Japan, Switzerland, the Netherlands, Belgium, the United Kingdom, and the United States, all of which have highly liquid options markets and readily available model-free implied variances for their respective aggregate market indexes (see Siriopoulos and Fassas (2009) for a recent summary of the model-free and parametric option-implied volatility indexes available for different countries). We begin with a brief discussion of the data.

A. Data and Summary Statistics

Our monthly aggregate market returns for the different countries are based on data for the French CAC 40 (obtained from Euronext), the German DAX 30 (obtained from Deutsche Börse), the Japanese Nikkei 225, the Swiss SMI 20, the Netherlands AEX, the Belgium BEL 20, and the U.K. FTSE 100 (all obtained from Datastream), and the U.S. S&P 500 (obtained from S&P). We use the sum of the daily squared returns over a month to construct end-of-month realized variances RV_t^i for each of the countries. We obtained the corresponding end-of-month model-free implied volatilities $(IV_t^i)^{1/2}$ for the S&P 500 (VIX) from the CBOE, the CAC (VCAC) from Euronext, and the DAX (VDAX) from Deutsche Börse, while those for the FTSE (VFTSE), SMI (VSMI), AEX (VAEX), and BEL (VBEL) were obtained from Datastream. Our data for the Japanese volatility index (VXJ) were obtained directly from the Center for the Study of Finance and Insurance at Osaka University (see Nishina, Maghrebi, and Kim (2006) for a more detailed discussion of the VXJ index). Finally, the risk-free rates used in the construction of the excess returns were obtained from the Federal Reserve Board and Eurocurrency via Datastream.¹¹

The sample period for each of the series extends from Jan. 2000 to Dec. 2011. The beginning of the sample coincides with the back-dated initial date of the NYSE Euronext volatility indices.¹² The use of more recent data through 2011

¹¹The use of excess returns, as opposed to raw returns, has almost no effect on the results from the return predictability regressions we report.

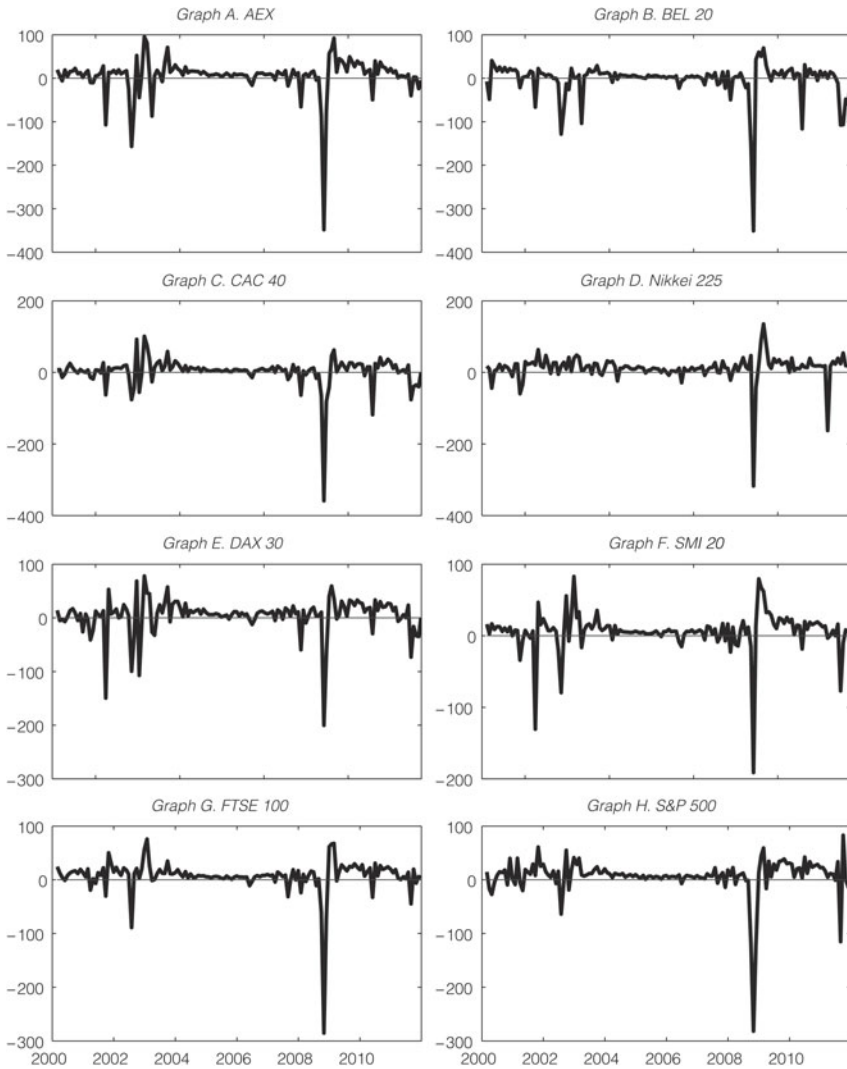
¹²The volatility indexes are available prior to Jan. 2000 for some of the countries: VDAX (Dec. 1994), VXJ (Jan. 1998), VSMI (Jan. 1999), and VIX (Jan. 1990). Comparable results to those for the country-specific regressions discussed below based on the longest possible sample for each of the countries are reported in a Supplementary Appendix available from the authors.

allows for additional validation of the original empirical evidence for the United States based on data prior to the financial crisis.

In accordance with the empirical analysis in the previous section, the proxy for the variance risk premium for each of the individual countries is simply defined by $VRP_t^i \equiv IV_t^i - RV_{t-1,t}^i$. As we note, this proxy has the obvious advantage of being directly observable. The time-series plots of VRP_t^i for each of the eight countries in Figure 3 clearly show the dramatic impact of the financial

FIGURE 3
Variance Risk Premia

Figure 3 shows the monthly proxies for the variance risk premia VRP_t^i for the Netherlands (AEX), Belgium (BEL 20), France (CAC 40), Japan (Nikkei 225), Germany (DAX 30), Switzerland (SMI 20), the United Kingdom (FTSE 100), and the United States (S&P 500). The risk premia are constructed by subtracting the actual realized variation from the model-free option-implied variation. The sample period spans Jan. 2000 to Dec. 2011.



crisis and the exceptionally large variance risk premia observed in the Fall of 2008. Interestingly, however, the premium for the DAX 30, and to a lesser extent the SMI 20, was almost as large and negative in 2001–2002.

The standard set of summary statistics reported in Table 2 also shows a remarkable coherence in the distributions of the variance risk premia and monthly excess returns across countries. In particular, looking at Panel A, the average excess returns all reflect the often-called “lost decade,” ranging from a high of -2.54 for Switzerland to a low of -9.26 for Belgium. Of course, the corresponding standard deviations all point to considerable variations in the returns around their negative sample means.

TABLE 2
Summary Statistics

Table 2 presents the monthly excess returns in annualized percentage form. The variance risk premia are in monthly percentage-squared form. The global index of variance risk premium is defined in the main text. The sample period extends from Jan. 2000 to Dec. 2011.

Panel A. Excess Returns and Variance Risk Premia

	AEX		BEL 20		CAC 40		DAX 30		FTSE 100		Nikkei 225		SMI 20		S&P 500		Global Index
	$r_t - r_{f,t}$	VRP _t	$r_t - r_{f,t}$	VRP _t	$r_t - r_{f,t}$	VRP _t	$r_t - r_{f,t}$	VRP _t	$r_t - r_{f,t}$	VRP _t	$r_t - r_{f,t}$	VRP _t	$r_t - r_{f,t}$	VRP _t	$r_t - r_{f,t}$	VRP _t	
Mean	-9.26	6.37	-5.25	-2.74	-8.64	2.75	-5.00	4.68	-4.93	7.55	-7.59	12.32	-2.54	6.45	-2.94	7.40	7.38
Std. dev.	76.59	43.30	65.02	43.28	67.49	41.69	82.81	33.43	52.35	32.16	71.73	38.47	51.90	28.10	57.52	35.47	32.74
Skewness	-0.99	-4.47	-1.36	-4.58	-0.59	-4.84	-0.90	-2.83	-0.64	-5.56	-0.76	-4.88	-0.70	-3.25	-0.61	-4.71	-5.46
Kurtosis	5.01	35.00	6.44	32.81	3.60	41.98	5.49	15.98	3.56	50.72	4.66	42.71	3.50	23.86	3.92	35.55	47.52
AR(1)	0.12	0.35	0.30	0.31	0.13	0.30	0.09	0.10	0.07	0.34	0.13	0.16	0.27	0.14	0.15	0.39	0.36

Panel B. Correlation Matrix for Excess Returns

	AEX	BEL 20	CAC 40	DAX 30	FTSE 100	Nikkei 225	SMI 20	S&P 500
AEX	1.00	0.86	0.91	0.88	0.87	0.61	0.83	0.81
BEL 20		1.00	0.83	0.76	0.81	0.52	0.79	0.75
CAC 40			1.00	0.93	0.90	0.60	0.84	0.87
DAX 30				1.00	0.85	0.57	0.80	0.83
FTSE 100					1.00	0.62	0.80	0.88
Nikkei 225						1.00	0.58	0.64
SMI 20							1.00	0.78
S&P 500								1.00

Panel C. Correlation Matrix for Variance Risk Premia

	AEX	BEL 20	CAC 40	DAX 30	FTSE 100	Nikkei 225	SMI 20	S&P 500	Global
AEX	1.00	0.85	0.91	0.86	0.92	0.64	0.86	0.81	0.85
BEL 20		1.00	0.81	0.68	0.80	0.51	0.77	0.69	0.93
CAC 40			1.00	0.84	0.89	0.65	0.79	0.82	0.86
DAX 30				1.00	0.78	0.54	0.86	0.70	0.92
FTSE 100					1.00	0.73	0.84	0.88	0.64
Nikkei 225						1.00	0.60	0.64	0.87
SMI 20							1.00	0.69	0.81
S&P 500								1.00	0.89
Global									1.00

The variance risk premia are almost all positive, on average, ranging from a low of -2.74 for Belgium to a high of 12.32 for Japan on a percentage-squared monthly basis. “Selling” volatility has been highly profitable, on average, over the last decade. Meanwhile, consistent with the visual impressions from Figure 3,

all of the premia are significantly negatively skewed and exhibit large excess kurtosis. Even though the implied and realized variances are both strongly serially correlated for all of the countries, the variance risk premia are generally not very persistent, and the maximum first-order serial correlation observed for the S&P 500 equals just 0.39. Turning to Panels B and C of Table 2, the sample cross-country correlations are all fairly high, and with the exception of those for the Nikkei and Belgium, the correlations for the returns all exceed 0.80, while those for the variance risk premia are in excess of 0.70.

The similarities in the summary statistics in Table 2 and the time-series plots in Figure 3 naturally suggest that the same predictive relationship documented for the U.S. returns and variance risk premium may hold true for the other countries. The results discussed in the next subsection generally corroborate this conjecture.

B. Country-Specific Regressions

In parallel to the general multiperiod return regressions defined in equation (1), our monthly return regressions for each of the individual countries may be conveniently expressed as

$$(2) \quad h^{-1}r_{t,t+h}^i = a^i(h) + b^i(h)\text{VRP}_t^i + u_{t,t+h}^i,$$

where $r_{t,t+h}^i$ and VRP_t^i refer to the $h = 1, 2, \dots, 12$ month excess return and variance risk premium for country i , respectively.

The actual estimates for $b^i(h)$ and the corresponding t^{NW} -statistics reported in Table 3 obviously differ somewhat across countries. However, with the exception of France, Belgium, and the United States, the estimated coefficients all show the same general pattern, starting out fairly low and insignificant at the shortest 1-month horizon, rising to their largest values at 3–5 months, and then gradually tapering off thereafter for longer return horizons. These similarities are also evident in Figure 4, which displays the regression coefficients along with their 90% NW (1987) standard error bands according to our simulated critical value in the simulation section.¹³

These similarities in the patterns in the estimated $b(h)$ coefficients naturally translate into very similar patterns in the regression $R^2(h)$ s as well. In particular, looking at the plots in Figure 5, all of the $R^2(h)$ s exhibit an almost identical hump-shaped pattern with the degree of predictability maximized around the 4-month horizon. Of course, the actual values of the $R^2(h)$ s vary somewhat across

¹³We also use the Stambaugh (1999) correction for the country-specific regression, and we find that the estimated bias is negligible. In fact, variance risk premia at monthly frequencies are much less persistent, and the contemporaneous correlations between residuals of bivariate VARs are only slightly negatively correlated. The Stambaugh correction results are reported in the Supplementary Appendix (available from the authors).

TABLE 3
Country-Specific Regressions

Table 3 presents the results based on the monthly regression in equation (2). t^{NW} statistics are reported in parentheses. The sample period extends from Jan. 2000 to Dec. 2011.

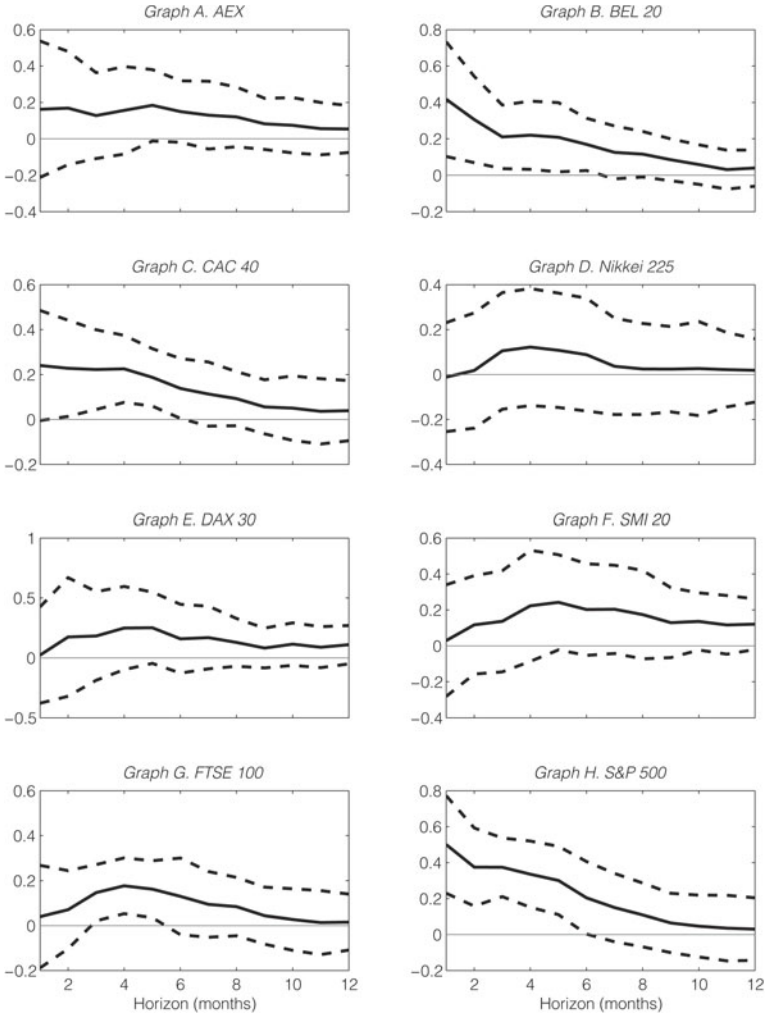
Index		Horizon							
		1	2	3	4	5	6	9	12
AEX	Constant	-10.29 (-1.37)	-10.55 (-1.44)	-10.31 (-1.42)	-10.73 (-1.49)	-11.10 (-1.55)	-10.79 (-1.49)	-10.15 (-1.36)	-9.74 (-1.29)
	VRP_t^j	0.16 (0.98)	0.17 (1.37)	0.13 (1.51)	0.16 (1.91)	0.18 (3.04)	0.15 (2.86)	0.08 (1.94)	0.05 (1.47)
	Adj. R^2	0.14	0.91	0.58	1.69	3.20	2.25	0.50	-0.07
BEL 20	Constant	-4.11 (-0.67)	-4.58 (-0.69)	-4.85 (-0.68)	-5.04 (-0.69)	-5.20 (-0.70)	-5.20 (-0.69)	-4.93 (-0.63)	-4.79 (-0.61)
	VRP_t^j	0.42 (3.00)	0.31 (3.25)	0.21 (3.36)	0.22 (3.44)	0.21 (3.56)	0.17 (3.79)	0.08 (2.45)	0.04 (1.38)
	Adj. R^2	7.07	5.59	3.35	4.53	4.33	2.93	0.47	-0.44
CAC 40	Constant	-9.30 (-1.55)	-9.38 (-1.56)	-9.46 (-1.56)	-9.77 (-1.58)	-9.75 (-1.54)	-9.40 (-1.46)	-8.60 (-1.28)	-8.20 (-1.19)
	VRP_t^j	0.24 (2.22)	0.23 (2.68)	0.22 (3.49)	0.23 (4.45)	0.19 (4.75)	0.14 (3.33)	0.06 (1.54)	0.04 (1.03)
	Adj. R^2	1.50	2.78	4.15	5.55	4.17	2.32	-0.05	-0.33
DAX 30	Constant	-5.11 (-0.65)	-5.73 (-0.75)	-5.70 (-0.76)	-6.18 (-0.84)	-6.27 (-0.85)	-5.41 (-0.72)	-4.35 (-0.56)	-4.14 (-0.53)
	VRP_t^j	0.02 (0.13)	0.17 (0.89)	0.18 (1.38)	0.25 (2.11)	0.25 (2.74)	0.16 (1.79)	0.08 (1.64)	0.11 (2.41)
	Adj. R^2	-0.71	0.19	0.75	2.70	3.28	1.12	-0.08	0.83
FTSE 100	Constant	-5.23 (-1.08)	-5.71 (-1.20)	-6.29 (-1.38)	-6.71 (-1.46)	-6.69 (-1.42)	-6.37 (-1.32)	-5.60 (-1.12)	-5.33 (-1.04)
	VRP_t^j	0.04 (0.39)	0.07 (1.03)	0.15 (3.26)	0.18 (4.19)	0.16 (4.14)	0.13 (2.45)	0.04 (1.15)	0.02 (0.44)
	Adj. R^2	-0.65	-0.36	1.56	3.64	3.40	2.25	-0.30	-0.70
Nikkei 225	Constant	-7.45 (-1.11)	-7.94 (-1.22)	-8.55 (-1.31)	-8.38 (-1.32)	-8.07 (-1.27)	-7.46 (-1.17)	-6.16 (-0.93)	-5.57 (-0.83)
	VRP_t^j	-0.01 (-0.11)	0.02 (0.18)	0.11 (1.13)	0.12 (1.38)	0.11 (1.38)	0.09 (1.13)	0.02 (0.41)	0.02 (0.46)
	Adj. R^2	-0.71	-0.70	0.08	0.60	0.50	0.24	-0.65	-0.70
SMI 20	Constant	-2.73 (-0.46)	-3.78 (-0.65)	-4.01 (-0.69)	-4.84 (-0.85)	-5.27 (-0.93)	-5.07 (-0.89)	-4.45 (-0.76)	-4.36 (-0.72)
	VRP_t^j	0.03 (0.22)	0.12 (1.08)	0.14 (1.35)	0.22 (2.12)	0.24 (2.97)	0.20 (2.56)	0.13 (2.20)	0.12 (3.01)
	Adj. R^2	-0.69	-0.08	0.42	2.98	4.11	3.08	1.36	1.62
S&P 500	Constant	-6.64 (-1.46)	-6.25 (-1.33)	-6.34 (-1.36)	-6.09 (-1.26)	-6.17 (-1.24)	-5.31 (-1.04)	-4.12 (-0.77)	-3.68 (-0.68)
	VRP_t^j	0.50 (4.17)	0.38 (4.36)	0.37 (6.39)	0.34 (5.37)	0.30 (5.13)	0.20 (3.26)	0.06 (1.30)	0.03 (0.61)
	Adj. R^2	8.89	8.72	13.03	12.83	10.77	5.26	0.10	-0.53

the different country indices, achieving a maximum of only 0.60% for the Nikkei 225 compared to 13.03% for the S&P 500.¹⁴ Interestingly, this value of adjusted $R^2(3) = 13.03\%$ for the United States exceeds that obtained with monthly data

¹⁴This lack of predictability for Japan is also consistent with the evidence reported in Ubukata and Watanabe (2011).

FIGURE 4
Country-Specific Regression Coefficients

Figure 4 shows the estimated regression coefficients for VRP_t^i for each of the country-specific return regressions reported in Table 3, together with NW (1987) based 90% standard error bands; see Table 1 for simulated critical values from 1 to 12 months. The regressions are based on monthly data from Jan. 2000 to Dec. 2011.

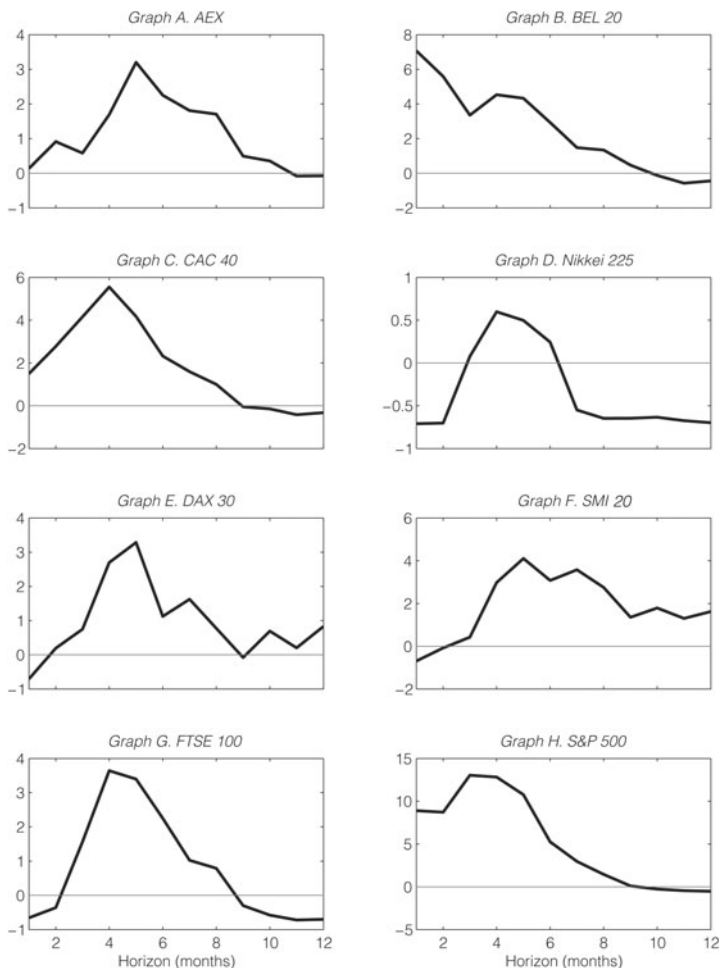


through the end of 2007 previously reported in BTZ (2009) and Drechsler and Yaron (2011).

The qualitative results from the country-specific VRP regressions, while not as significant, are generally in line with the existing results for the United States. Going one step further, the similarities in the patterns observed across the different countries also suggest that even stronger results may be available by pooling the regressions and entertaining the notion of a common global variance risk premium. We explore these ideas next.

FIGURE 5
Country-Specific Regression R^2 s

Figure 5 shows the adjusted $R^2(h)$ s for the country-specific return regressions reported in Table 3. The regressions are based on monthly data from Jan. 2000 to Dec. 2011.



IV. Global Variance Risk

Our proxy for the global variance risk premium is based on a simple capitalization weighted average of the proxies for country-specific variance risk premia,

$$VRP_t^{GLOBAL} \equiv \sum_{i=1}^8 w_t^i VRP_t^i,$$

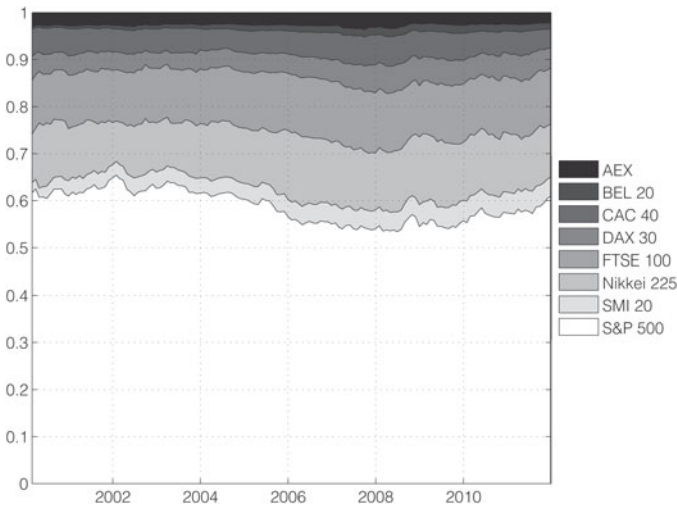
where $i = 1, 2, \dots, 8$ refers to each of the eight countries included in our analysis.¹⁵ The end-of-month market capitalizations used in defining the weights

¹⁵This parallels the construction used in Harvey (1991) in the estimation of the world price of covariance risk.

w_t^i are obtained from Thomson Reuters Institutional Brokers' Estimate System (IBES) via Datastream. The plot of the weights in Figure 6 shows that the U.S. market accounts for around 60% through most of the sample period, with the Japanese market a distant second. This large weight assigned to the U.S. market in our definition of the global VRP index is also implicit in the aforementioned summary statistics in Panel C in Table 2 and the relatively high correlation of 0.89 between VRP_t^{GLOBAL} and VRP_t^{SP500} .

FIGURE 6
Market Capitalization

Figure 6 shows the relative market capitalization by aggregate index for the Netherlands (AEX), Belgium (BEL 20), France (CAC 40), Germany (DAX 30), the United Kingdom (FTSE 100), Japan (Nikkei 225), Switzerland (SMI 20), and the United States (S&P 500).



A. Individual Country Regressions

The results for the regressions obtained by replacing the country-specific VRP_t^i s in equation (2) with the new VRP_t^{GLOBAL} proxy,

$$(3) \quad h^{-1}r_{t,t+h}^i = a^i(h) + b^i(h)VRP_t^{GLOBAL} + u_{t,t+h}^i,$$

are reported in Table 4. Comparing the results to those for the country-specific regressions in Table 3 reveals even stronger commonalities and uniform patterns across countries. The global VRP proxy serves as a highly significant predictor variable for all of the different country returns, with t^{NW} -statistics systematically in excess of 4.0 at the 4- or 5-month horizon. Further increasing the horizon h , VRP_t^{GLOBAL} systematically becomes insignificant for predicting the longer 9- and 12-month returns.

These striking cross-country similarities are also evident from the plots of the estimated regression coefficients and the 90% NW (1987) based confidence bands with simulated critical values in Figure 7. Not only do the individual country estimates for the $b^i(h)$ s look very similar, the confidence bands also tend to

TABLE 4
Global Variance Risk Premium Regressions

Table 4 presents the results based on the monthly regression in equation (3). t^{NW} -statistics are reported in parentheses. The sample period extends from Jan. 2000 to Dec. 2011.

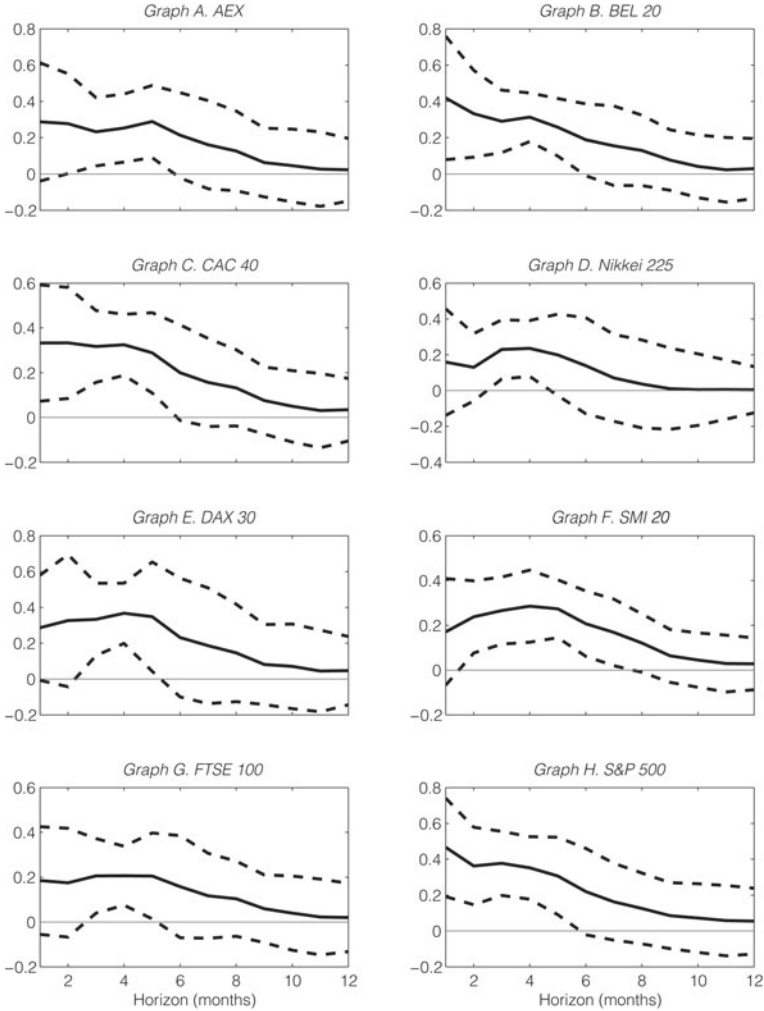
Index		Horizon							
		1	2	3	4	5	6	9	12
AEX	Constant	-11.38 (-1.56)	-11.52 (-1.59)	-11.22 (-1.58)	-11.52 (-1.61)	-12.09 (-1.68)	-11.39 (-1.56)	-10.06 (-1.34)	-9.55 (-1.26)
	VRP_t^{GLOBAL}	0.29 (1.99)	0.28 (2.54)	0.23 (3.45)	0.25 (3.98)	0.29 (4.72)	0.21 (2.92)	0.06 (1.11)	0.02 (0.48)
	Adj. R^2	0.80	1.80	1.77	2.86	4.49	2.51	-0.36	-0.70
BEL 20	Constant	-8.35 (-1.18)	-7.78 (-1.09)	-7.47 (-1.04)	-7.58 (-1.05)	-7.33 (-0.99)	-6.72 (-0.89)	-5.58 (-0.72)	-5.05 (-0.65)
	VRP_t^{GLOBAL}	0.42 (2.79)	0.33 (3.50)	0.29 (4.70)	0.31 (6.87)	0.26 (5.26)	0.19 (3.06)	0.08 (1.53)	0.03 (0.63)
	Adj. R^2	3.77	3.62	3.81	5.67	3.83	1.95	-0.16	-0.66
CAC 40	Constant	-11.09 (-1.81)	-11.18 (-1.84)	-11.12 (-1.85)	-11.29 (-1.85)	-11.24 (-1.77)	-10.38 (-1.61)	-8.96 (-1.34)	-8.31 (-1.21)
	VRP_t^{GLOBAL}	0.33 (2.89)	0.33 (3.38)	0.32 (5.51)	0.32 (6.99)	0.29 (5.22)	0.20 (3.01)	0.08 (1.67)	0.03 (0.87)
	Adj. R^2	1.90	3.95	5.51	7.34	6.30	3.12	0.01	-0.57
DAX 30	Constant	-7.12 (-0.96)	-7.31 (-0.98)	-7.26 (-1.01)	-7.49 (-1.03)	-7.52 (-1.00)	-6.27 (-0.82)	-4.49 (-0.58)	-3.86 (-0.49)
	VRP_t^{GLOBAL}	0.29 (2.21)	0.33 (2.23)	0.33 (4.59)	0.37 (6.41)	0.35 (3.70)	0.23 (2.25)	0.08 (1.20)	0.05 (0.86)
	Adj. R^2	0.58	2.36	4.05	6.47	6.61	2.98	-0.13	-0.50
FTSE 100	Constant	-6.30 (-1.41)	-6.50 (-1.44)	-6.73 (-1.53)	-6.87 (-1.52)	-7.03 (-1.51)	-6.57 (-1.37)	-5.71 (-1.15)	-5.37 (-1.05)
	VRP_t^{GLOBAL}	0.18 (1.74)	0.18 (1.82)	0.21 (3.45)	0.21 (4.63)	0.21 (3.48)	0.16 (2.23)	0.06 (1.30)	0.02 (0.47)
	Adj. R^2	0.63	1.56	3.97	5.33	5.80	3.57	0.04	-0.65
Nikkei 225	Constant	-8.76 (-1.25)	-8.69 (-1.29)	-9.00 (-1.41)	-8.62 (-1.37)	-8.36 (-1.31)	-7.51 (-1.17)	-5.97 (-0.91)	-5.37 (-0.81)
	VRP_t^{GLOBAL}	0.16 (1.21)	0.13 (1.74)	0.23 (3.88)	0.24 (4.45)	0.20 (2.84)	0.14 (1.67)	0.01 (0.15)	0.00 (0.13)
	Adj. R^2	-0.18	-0.09	2.03	2.81	2.12	0.91	-0.74	-0.77
SMI 20	Constant	-3.80 (-0.67)	-4.80 (-0.87)	-5.14 (-0.94)	-5.44 (-0.98)	-5.69 (-1.01)	-5.23 (-0.91)	-4.01 (-0.68)	-3.71 (-0.61)
	VRP_t^{GLOBAL}	0.17 (1.62)	0.24 (3.71)	0.27 (4.94)	0.29 (5.21)	0.27 (6.86)	0.21 (4.56)	0.06 (1.78)	0.03 (0.85)
	Adj. R^2	0.45	2.86	5.16	7.38	7.60	4.65	-0.07	-0.60
S&P 500	Constant	-6.38 (-1.39)	-6.13 (-1.30)	-6.34 (-1.37)	-6.26 (-1.32)	-6.20 (-1.27)	-5.42 (-1.08)	-4.28 (-0.80)	-3.87 (-0.71)
	VRP_t^{GLOBAL}	0.47 (3.84)	0.36 (4.22)	0.38 (5.89)	0.35 (5.94)	0.31 (4.60)	0.22 (2.94)	0.09 (1.54)	0.05 (1.04)
	Adj. R^2	6.32	6.74	11.10	12.05	9.94	5.47	0.57	-0.06

be tighter compared to the country-specific regressions discussed above. Furthermore, along these lines, Figure 8 shows the general patterns in the predictability, as measured by the $R^2(h)$ s, to be very similarly shaped across countries, with peaks at the 4- to 5-month return horizon.¹⁶

¹⁶The relatively large weight assigned to the United States in our construction of the global variance risk premium means that fairly similar results are obtained by replacing the new VRP_t^{GLOBAL} in the regressions in equation (3) with VRP_t^{SP500} . These additional results are available from the authors. Comparable empirical results based on the U.S. variance risk premium have also recently been

FIGURE 7
Global VRP Regression Coefficients

Figure 7 shows the coefficient estimates for VRP_t^{GLOBAL} from the return regressions reported in Table 4, together with NW (1987) based 90% standard error bands; see Table 1 for the simulated critical values from 1 to 12 months. The regressions are based on monthly data from Jan. 2000 to Dec. 2011.

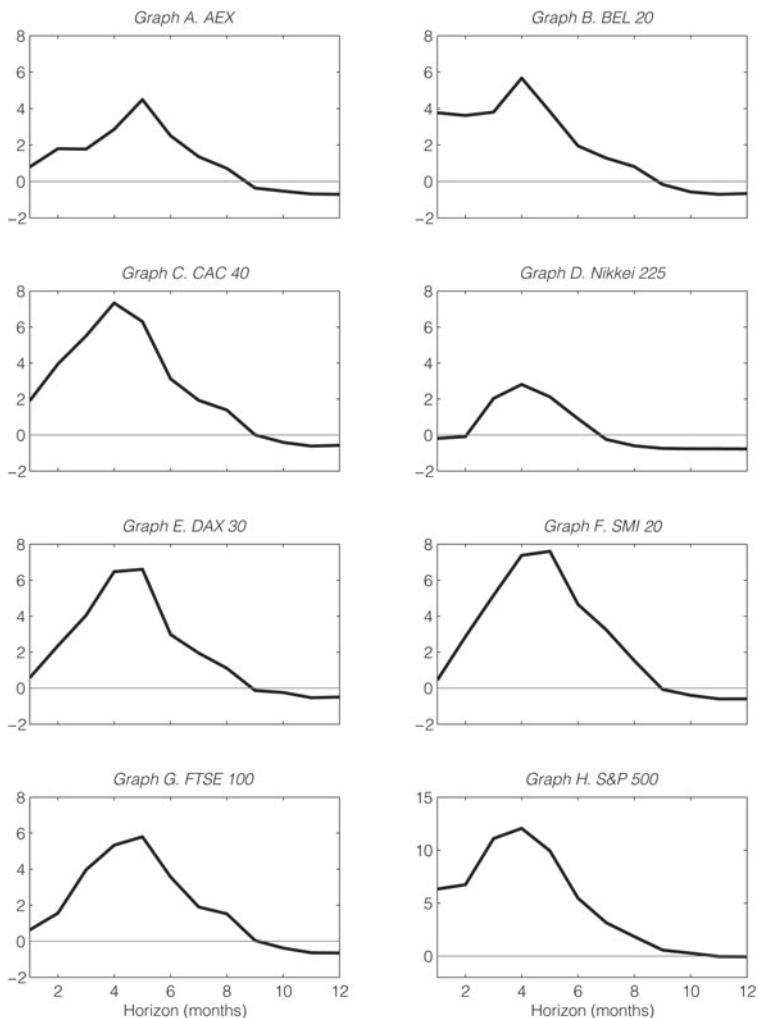


These remarkable similarities in the estimates for the different countries naturally suggest restricting the coefficients in equation (3) to be the same across countries, as a way to enhance the efficiency of the estimates and to ensure a common reward for bearing global variance risk.

reported in concurrent independent work by Londono (2011), who ascribes the predictability to informational frictions along the lines of Rapach et al. (2013).

FIGURE 8
Global VRP Regression R^2 s

Figure 8 shows the adjusted $R^2(h)$ s from regressing the individual country returns on the VRP_t^{GLOBAL} reported in Table 4. The regressions are based on monthly data from Jan. 2000 to Dec. 2011.



B. Panel Regressions

The estimation results from the panel regression that restrict the coefficients for the global variance risk premium to be the same across countries,

$$(4) \quad h^{-1}r_{t,t+h}^i = a(h) + b(h)VRP_t^{GLOBAL} + u_{t,t+h}^i,$$

are reported in Table 5 (for additional details on calculating standard errors, see, e.g., Petersen (2009)).¹⁷ As the table clearly shows, the use of panel

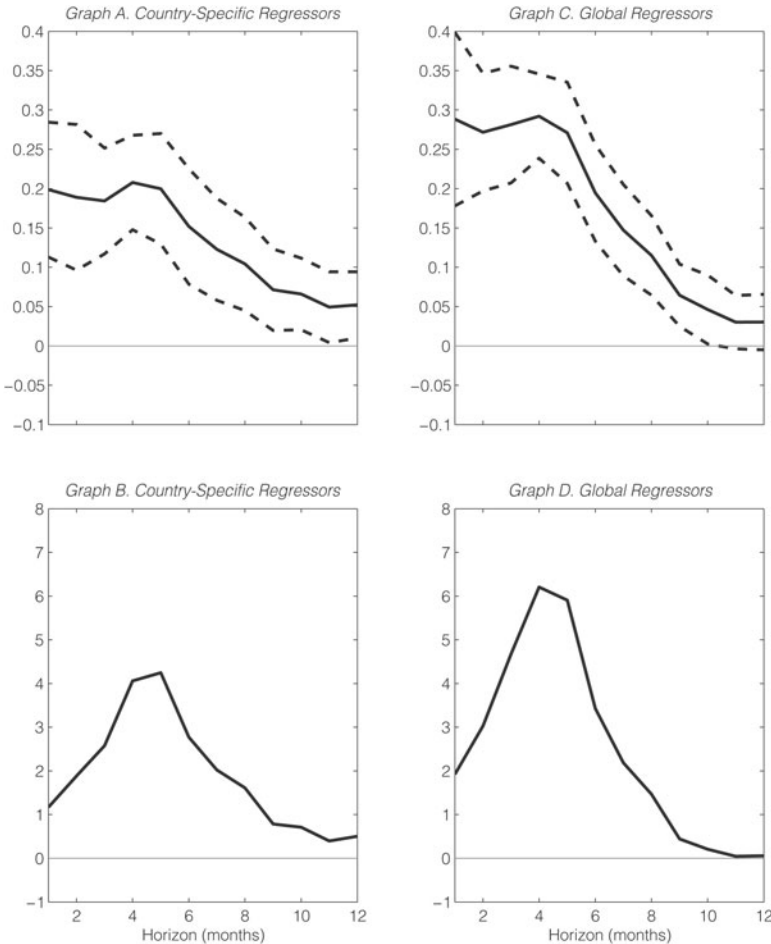
¹⁷We also experimented with the two-way cluster analysis in Cameron, Gelbach, and Miller (2011), resulting in very similar findings.

regressions does indeed result in more accurate estimates and highly significant t^{NW} -statistics of 10.91 at the 4-month horizon. The average panel regression $R^2(h)$ s for the eight countries also gradually rise from less than 2% at the 1-month horizon to a large 6.21% for the 4-month returns, tapering off to 0 for the longer 9- to 12-month return horizons.

These key empirical findings are succinctly summarized in Figure 9, which plots the panel regression estimates for the $b(h)$ s based on the country-specific VRPs and the global VRP proxy along with their two NW (1987) based standard error bands (Graphs A and C) and the corresponding panel regression $R^2(h)$ s

FIGURE 9
Panel Regression Coefficients and R^2 s

Graphs A and C of Figure 9 show the estimated panel regression coefficients from regressing the returns on the individual country variance risk premia VRP_i^c and the global variance risk premia $\text{VRP}_i^{\text{GLOBAL}}$, respectively, reported in Table 5, together with two NW (1987) based standard error bands. Graphs B and D show the $R^2(h)$ s from the same two panel regressions. The regressions are based on monthly data from Jan. 2000 to Dec. 2011.

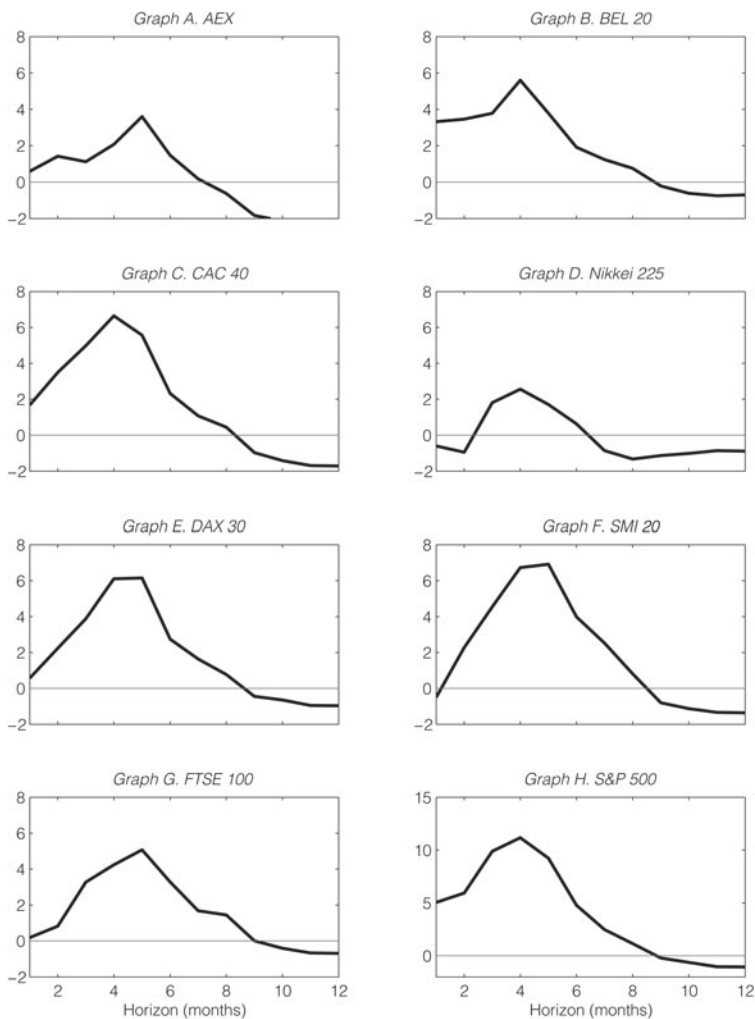


(Graphs B and D). The VRP^{GLOBAL} -based regressions (depicted in Graphs C and D) obviously result in sharper coefficient estimates and stronger average predictability across the eight countries than do the individual country VRP^i regressions (depicted in Graphs A and B).

The average panel regression $R^2(h)$ s, of course, mask important cross-country differences in the degree of predictability. We therefore also show in Figure 10 the country-specific implied $R^2(h)$ s obtained by evaluating the individual country regressions in equation (3) at the more precisely estimated common $\hat{a}(h)$ and $\hat{b}(h)$ obtained from the panel regressions in equation (4).

FIGURE 10
Global VRP Panel Regression R^2 s

Figure 10 shows the adjusted $R^2(h)$ s implied by the VRP^{GLOBAL} panel regressions reported in Panel A Table 5. The regressions are based on monthly data from Jan. 2000 to Dec. 2011.



Comparing Figure 10 to Figure 8 for the individual country regressions, it is clear that the added precision afforded by restricting the $a^i(h)$ and $b^i(h)$ coefficients to be the same across countries sacrifices very little in terms of the implied predictability.

TABLE 5
Panel Regressions

Table 5 presents the results based on the monthly global and country-specific panel regressions in equations (4) and (5), respectively. NW (1987) based *t*-statistics are reported in the parentheses. The sample period extends from Jan. 2000 to Dec. 2011.

	Horizon							
	1	2	3	4	5	6	9	12
<i>Panel A. Global Regressors</i>								
Constant	-7.90 (-3.08)	-7.99 (-3.98)	-8.04 (-4.56)	-8.13 (-5.06)	-8.18 (-5.29)	-7.44 (-4.88)	-6.13 (-4.63)	-5.64 (-5.89)
VRP _t ^{GLOBAL}	0.29 (5.23)	0.27 (7.27)	0.28 (7.58)	0.29 (10.91)	0.27 (8.44)	0.19 (6.33)	0.06 (3.26)	0.03 (1.72)
Adj. R ²	1.92	3.03	4.67	6.21	5.91	3.43	0.44	0.06
Constant	3.09 (0.14)	-1.20 (-0.06)	0.17 (0.01)	7.60 (0.56)	9.71 (0.70)	9.67 (0.72)	6.96 (0.72)	-1.59 (-0.22)
VRP _t ^{GLOBAL}	0.30 (5.95)	0.28 (9.41)	0.29 (9.97)	0.31 (13.58)	0.28 (10.79)	0.21 (9.04)	0.07 (5.19)	0.03 (2.30)
log(P _t ⁱ /E _t ⁱ) ^{GLOBAL}	-5.07 (-0.49)	-3.13 (-0.34)	-3.78 (-0.52)	-7.24 (-1.14)	-8.23 (-1.28)	-7.86 (-1.25)	-5.99 (-1.32)	-1.85 (-0.58)
Adj. R ²	1.88	2.97	4.65	6.38	6.21	3.74	0.66	0.00
<i>Panel B. Country-Specific Regressors</i>								
Constant	-6.87 (-2.56)	-7.03 (-2.72)	-6.98 (-3.22)	-7.24 (-3.52)	-7.32 (-4.09)	-6.89 (-4.29)	-6.08 (-5.11)	-5.73 (-5.27)
VRP _t ⁱ	0.20 (4.64)	0.19 (4.08)	0.18 (5.48)	0.21 (6.92)	0.20 (5.69)	0.15 (4.14)	0.07 (2.76)	0.05 (2.47)
Adj. R ²	1.17	1.88	2.57	4.06	4.25	2.77	0.79	0.50
Constant	2.31 (0.40)	3.15 (0.57)	2.80 (0.54)	2.97 (0.68)	3.15 (0.81)	2.93 (0.78)	2.48 (0.87)	0.33 (0.15)
VRP _t ⁱ	0.21 (5.58)	0.20 (5.62)	0.19 (9.51)	0.22 (10.47)	0.21 (7.24)	0.16 (5.39)	0.08 (3.78)	0.06 (3.61)
log(P _t ⁱ /E _t ⁱ) ⁱ	-4.61 (-1.79)	-5.11 (-2.22)	-4.91 (-2.24)	-5.12 (-2.63)	-5.24 (-2.98)	-4.91 (-2.82)	-4.27 (-3.05)	-3.02 (-2.62)
Adj. R ²	1.18	2.00	2.75	4.33	4.60	3.12	1.15	0.71

C. Robustness Checks

To assess the robustness of these striking international predictability patterns, Panel B of Table 5 reports the results obtained by including a capitalization weighted average of the country-specific PE ratios as an additional regressor in equation (4). Consistent with the results for the U.S. market in isolation reported in BTZ (2009), the global PE ratio adds nothing to the predictability afforded by VRP^{GLOBAL} within the 1-year horizons reported in the table, leaving all of the estimates for $b(h)$ and the $R^2(h)$ s almost the same. The predictability of the global variance risk premium is effectively “orthogonal” to that documented in the existing literature based on more traditional macrofinance variables, such as the PE ratio, dividend yields, and consumption-wealth ratios, which are typically

significant only over longer multiyear return horizons (see, e.g., the classic studies by Fama and French (1988), Campbell and Shiller (1988), and Lettau and Ludvigson (2001)).¹⁸

To further highlight the predictive gains afforded by the use of our global VRP as opposed to the own-country VRPs, Panel B of Table 5 shows the estimates obtained by including each individual country's premium in a panel regression in place of $\text{VRP}^{\text{GLOBAL}}$,

$$(5) \quad h^{-1}r_{t,t+h}^i = a(h) + b(h)\text{VRP}_t^i + u_{t,t+h}^i.$$

While the results still point to overall efficiency gains from the use of the panel regression relative to the country-specific regressions in Table 3, the magnitude of the return predictability is obviously much lower than for $\text{VRP}^{\text{GLOBAL}}$. The global variance risk premium is clearly a much better predictor of the future returns for most of the countries than the individual country-specific premia. Again, including the country-specific PE ratios in the same panel regression does not materially affect the overall predictability as measured by the R^2 's nor the values of the estimated regression coefficients for the variance risk premia.

D. Forward-Looking Global Variance Risk Premium

Our proxy for the global variance risk premium underlying our main findings discussed above is based on a weighted average of the variance difference for each of the countries. This directly mirrors the original proxy for the U.S. variance risk premium employed in BTZ (2009) and the proxy used in the country-specific regressions in Section III. To assess the sensitivity of our results to this simple and easy-to-implement proxy, we briefly summarize the results obtained by replacing the model-free lagged monthly realized variances with forward-looking model-based expectations in the way we define the global variance risk premium.

Specifically, let $E_t(\text{RV}_{t,t+1}^i)$ denote the time t expectation of the 1-month ahead return variation for country i . Additionally, let $\text{FVRP}_t^i = \text{IV}_t^i - E_t(\text{RV}_{t,t+1}^i)$ denote the corresponding forward-looking variance risk premia for country i . We then define a forward-looking global variance risk premium by

$$\text{FVRP}_t^{\text{GLOBAL}} \equiv \sum_{i=1}^8 w_t^i \text{FVRP}_t^i.$$

In contrast to the $\text{VRP}^{\text{GLOBAL}}$ defined above, $\text{FVRP}^{\text{GLOBAL}}$ necessitates the use of a model for generating the forward expectations $E_t(\text{RV}_{t,t+1}^i)$. In the results reported on below, we follow Andersen, Bollerslev, and Diebold (2007) and Corsi (2009) in generating these forecasts from heterogeneous autoregressive model of the realized volatility (HAR-RV) type models in which we regress $\text{RV}_{t,t+1}^i$ for

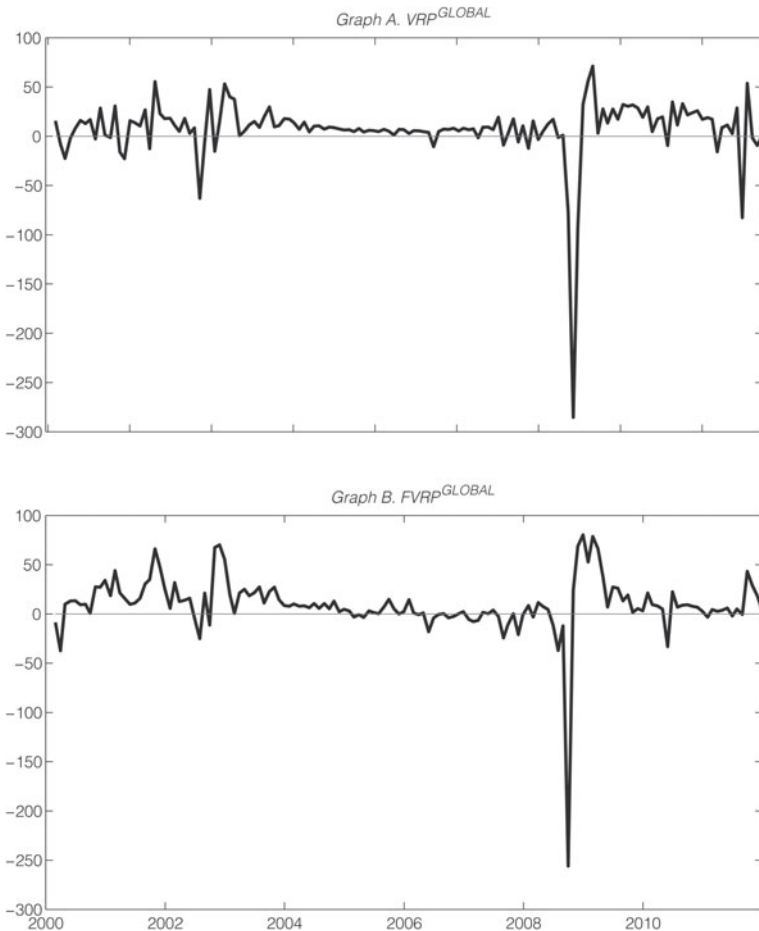
¹⁸Further corroborating the results for the U.S. market in BTZ (2009), we also found that including the implied global variance or the realized global variance together with the global variance risk premium resulted in mostly insignificant coefficient estimates. These additional results are available from the authors.

each of the eight countries on the daily, weekly, and monthly realized variances, $RV_{t-\Delta,t}^i$, $RV_{t-5\Delta,t}^i$, and $RV_{t-1,t}^i$ (where $\Delta = 1/20$), respectively, along with the option-implied variances, IV_t^i , for all of the other seven countries.¹⁹

The resulting $FVRP^{GLOBAL}$ is plotted in Figure 11 (Graph B), together with the previously used simple VRP^{GLOBAL} proxy (Graph A). While the two series

FIGURE 11
Variance Risk Premia

Figure 11 shows our proxies for the monthly global variance risk premia VRP_t^{GLOBAL} (Graph A) and $FVRP_t^{GLOBAL}$ (Graph B), as defined in the main text. The sample period spans Jan. 2000 to Dec. 2011.



¹⁹We make sure that all of the regressors are properly aligned to correct for the different time zones, so that none of the predictions involve any future information. We also experimented with the use of a standard VAR(1) model involving only the current monthly realized variation measures, $RV_{t-1,t}^i$, and option-implied variation measures, IV_t^i , for generating $E_t(RV_{t,t+1}^i)$, resulting in qualitatively similar, albeit not as significant, predictive return regressions. Further details concerning these additional results are summarized in the Supplementary Appendix available from the authors.

obviously differ, the general dynamic dependencies are obviously quite similar. The large negative spike in VRP^{GLOBAL} observed at the height of the financial crisis is slightly diminished in the forward-looking $FVRP^{GLOBAL}$.

Turning to the predictive return regressions, Panel A of Table 6 reports the estimates from the same panel regressions in equation (4) using $FVRP^{GLOBAL}$ in place of VRP^{GLOBAL} . While the NW (1987) based t -statistics for the 1- to 6-month returns are all slightly lower than the comparable t^{NW} -statistics reported in Panel A of Table 5, they remain highly significant at any reasonable level. In fact, the statistical significance of the regressions based on $FVRP^{GLOBAL}$ extends to at least the 9-month horizon. The R^2 s also show a similar hump-shaped pattern to those in Table 5 and Figure 9, with the predictability now maximized at the slightly longer 5- to 6-month horizon. This shift in the location of the peak is also consistent with the Monte Carlo results in Figure 2 and the slightly smaller first-order autocorrelation of 0.31 for $FVRP^{GLOBAL}$ compared to 0.36 for VRP^{GLOBAL} . Panel B of Table 6 again further corroborates our key empirical findings and the idea that the predictability inherent in the global variance risk premium is essentially orthogonal to that in the global PE ratio, which kicks in only over longer annual horizons.

TABLE 6
Panel Regressions with Forecasted Global Variance Risk Premium

Table 6 presents the results based on the monthly forecasted global panel regressions in equations (4) and (5), respectively. NW (1987) based t -statistics are reported in the parentheses. The sample period extends from Jan. 2000 to Dec. 2011.

	Horizon							
	1	2	3	4	5	6	9	12
<i>Panel A. Forecasted VRP</i>								
Constant	-10.43 (-3.50)	-9.24 (-3.58)	-8.51 (-3.56)	-8.54 (-3.90)	-9.01 (-4.42)	-8.70 (-4.61)	-7.20 (-4.72)	-6.39 (-5.37)
$FVRP_t^{GLOBAL}$	0.49 (6.05)	0.35 (6.26)	0.28 (5.19)	0.28 (5.98)	0.32 (6.19)	0.30 (5.07)	0.17 (3.75)	0.10 (2.49)
Adj. R^2	4.86	4.26	3.91	4.80	7.68	7.63	3.22	1.56
<i>Panel B. Forecasted VRP and PE</i>								
Constant	4.61 (0.19)	-2.52 (-0.12)	-3.14 (-0.18)	3.16 (0.21)	10.34 (0.73)	12.38 (1.02)	10.32 (1.03)	0.89 (0.13)
$FVRP_t^{GLOBAL}$	0.50 (6.79)	0.35 (7.47)	0.28 (7.01)	0.29 (8.02)	0.34 (7.87)	0.32 (6.24)	0.18 (5.46)	0.11 (3.42)
$\log(P_t/E_t)^{GLOBAL}$	-6.95 (-0.62)	-3.10 (-0.32)	-2.48 (-0.30)	-5.39 (-0.76)	-8.90 (-1.35)	-9.69 (-1.67)	-8.03 (-1.77)	-3.33 (-1.14)
Adj. R^2	4.86	4.21	3.85	4.86	8.05	8.15	3.70	1.59

In sum, the estimated regression coefficients for the global variance risk premium are fairly similar across countries, and with the exception of the United States, the $R^2(h)$ s for the panel regressions are generally larger for the global VRP than for the “local” VRPs.

These empirical findings are directly in line with a stylized two-country consumption-based equilibrium model. We show that the global variance risk premia include a relatively larger amount of the aggregate volatility uncertainty than the local variance risk premium from the smaller country (of which the

consumption weight is less than one-half, directly mirroring the global variance risk premium's stronger return predictability). Conversely, for the larger country (the United States), the local VRP gives rise to marginally higher slope coefficients than the global VRP.²⁰

V. Conclusion

A number of recent studies have argued that aggregate U.S. stock market return is predictable over relatively short 3- to 5-month horizons by the difference between option-implied and actual realized variation, or the so-called variance risk premium. We show that this newly documented predictability is not due to finite sample biases in the statistical inference procedures, and that the apparent hump-shaped pattern in the degree of predictability documented in overlapping monthly returns regressions is entirely consistent with the implications from an empirically realistic bivariate daily time-series model for the returns and variance risk premia.

Further corroborating the existing empirical evidence for the United States, we show that the same basic predictive relationship between future returns and current variance risk premia holds true with more recent out-of-sample data through 2011. We also show that the same basic results hold true for a set of seven other countries, although the magnitude of the predictability and the statistical significance of the own-country variance risk premium tend to be somewhat muted relative to those observed for the United States.

Meanwhile, employing a capitalization weighted global variance risk premium results in similarly shaped predictability patterns across return horizons for *all* of the countries in our sample, and uniformly larger *t*-statistics. Further restricting the regression coefficients and the compensation for global variance risk to be the same across countries, we find even stronger results and highly significant test statistics, with the degree of predictability maximized at the 4- to 5-month horizon.

The global variance risk premium may be seen as a proxy for worldwide aggregate economic uncertainty; therefore, global variance risk premia generally provide more accurate predictions of the future individual country returns than the own-country variance risk premia. Alternatively, the global variance risk premium may be interpreted as a measure of aggregate risk aversion (e.g., Bekaert, Engstrom, and Xing (2009)) or a summary measure of worldwide disagreements in beliefs (e.g., Buraschi et al. (2014)). All of these different economic mechanisms likely play some role in generating the strong international return predictability embodied in the global variance risk premium first documented here. We leave it for future research to sort out this important question.

References

- Andersen, T. G.; T. Bollerslev; and F. X. Diebold. "Roughing It Up: Including Jump Components in the Measurement, Modeling, and Forecasting of Return Volatility." *Review of Economics and Statistics*, 89 (2007), 701–720.

²⁰Details of the model setup and calibration evidence can be found in the online Supplementary Appendix (available from the authors).

- Ang, A., and G. Bekaert. "Stock Return Predictability: Is It There?" *Review of Financial Studies*, 20 (2007), 651–707.
- Baker, M.; R. Taliaferro; and J. Wurgler. "Predicting Returns with Managerial Decision Variables: Is There a Small-Sample Bias?" *Journal of Finance*, 61 (2006), 1711–1730.
- Bakshi, G.; G. Panayotov; and G. Skoulakis. "The Baltic Dry Index as a Predictor of Global Stock Returns, Commodity Returns, and Global Economic Activity." Working Paper, University of Maryland (2011).
- Bekaert, G.; E. Engstrom; and Y. Xing. "Risk, Uncertainty and Asset Prices." *Journal of Financial Economics*, 91 (2009), 59–82.
- Bekaert, G., and R. J. Hodrick. "Characterizing Predictable Components in Excess Returns on Equity and Foreign Exchange Markets." *Journal of Finance*, 47 (1992), 467–509.
- Bollerslev, T.; N. Sizova; and G. Tauchen. "Volatility in Equilibrium: Asymmetries and Dynamics Dependencies." *Review of Finance*, 16 (2012), 31–80.
- Bollerslev, T.; G. Tauchen; and H. Zhou. "Expected Stock Returns and Variance Risk Premia." *Review of Financial Studies*, 22 (2009), 4463–4492.
- Bollerslev, T., and J. M. Wooldridge. "Quasi-Maximum Likelihood Estimation and Inference in Dynamic Models with Time-Varying Covariances." *Econometric Reviews*, 11 (1992), 143–172.
- Boudoukh, J.; M. Richardson; and R. F. Whitelaw. "The Myth of Long-Horizon Predictability." *Review of Financial Studies*, 21 (2008), 1577–1605.
- Buraschi, A.; F. Trojani; and A. Vedolin. "Economic Uncertainty, Disagreement, and Credit Markets." *Management Science*, 60 (2014), 1281–1296.
- Cameron, A. C.; J. B. Gelbach; and D. L. Miller. "Robust Inference with Multiway Clustering." *Journal of Business and Economic Statistics*, 29 (2011), 238–249.
- Campbell, J. Y., and Y. Hamao. "Predicting Stock Returns in the United States and Japan: A Study of Long-Term Capital Market Integration." *Journal of Finance*, 47 (1992), 63–49.
- Campbell, J. Y.; A. W. Lo; and A. C. MacKinlay. *The Econometrics of Financial Markets*. Princeton, NJ: Princeton University Press (1997), 27–82.
- Campbell, J. Y., and R. J. Shiller. "Stock Prices, Earnings, and Expected Dividends." *Journal of Finance*, 43 (1988), 661–676.
- Campbell, J. Y., and M. Yogo. "Efficient Tests of Stock Return Predictability." *Journal of Financial Economics*, 81 (2006), 27–60.
- Corsi, F. "A Simple Approximate Long Memory Model of Realized Volatility." *Journal of Financial Econometrics*, 7 (2009), 174–196.
- Drechsler, I., and A. Yaron. "What's Vol Got to Do with It?" *Review of Financial Studies*, 24 (2011), 1–45.
- Engle, R. F. "Dynamic Conditional Correlation: A Simple Class of Multivariate GARCH Models." *Journal of Business and Economic Statistics*, 20 (2002), 339–350.
- Fama, E. F., and K. R. French. "Dividend Yields and Expected Stock Returns." *Journal of Financial Economics*, 22 (1988), 3–25.
- Ferson, W. E., and C. R. Harvey. "The Risk and Predictability of International Equity Returns." *Review of Financial Studies*, 6 (1993), 527–566.
- Ferson, W. E.; S. Sarkissian; and T. T. Simin. "Spurious Regressions in Financial Economics?" *Journal of Finance*, 58 (2003), 1393–1414.
- Gabaix, X. "Variable Rare Disasters: An Exactly Solved Framework for Ten Puzzles in Macro-Finance." *Quarterly Journal of Economics*, 127 (2012), 645–700.
- Goyal, A., and I. Welch. "A Comprehensive Look at the Empirical Performance of Equity Premium Prediction." *Review of Financial Studies*, 21 (2008), 1455–1508.
- Harvey, C. R. "The World Price of Covariance Risk." *Journal of Finance*, 46 (1991), 111–157.
- Hjalmarsson, E. "Predicting Global Stock Returns." *Journal of Financial and Quantitative Analysis*, 45 (2010), 49–80.
- Hodrick, R. J. "Dividend Yields and Expected Stock Returns: Alternative Procedures for Inference and Measurement." *Review of Financial Studies*, 5 (1992), 357–386.
- Kelly, B. T. "Tail Risk and Asset Prices." Working Paper, University of Chicago (2011).
- Lettau, M., and S. Ludvigson. "Consumption, Aggregate Wealth, and Expected Stock Returns." *Journal of Finance*, 56 (2001), 815–849.
- Londono, J.-M. "The Variance Risk Premium around the World." Working Paper, Federal Reserve Board, Washington, DC (2011).
- Mueller, P.; A. Vedolin; and H. Zhou. "Short-Run Bond Risk Premia." Working Paper, London School of Economics, and Federal Reserve Board, Washington, DC (2011).
- Newey, W. K., and K. D. West. "A Simple Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix." *Econometrica*, 55 (1987), 703–708.
- Nishina, K.; N. Maghrebi; and M.-S. Kim. "Stock Market Volatility and the Forecasting Accuracy of Implied Volatility Indices." Working Paper, Graduate School of Economics and School of International Public Policy, Osaka University, Japan (2006).

- Petersen, M. A. "Estimating Standard Errors in Finance Panel Data Sets: Comparing Approaches." *Review of Financial Studies*, 22 (2009), 435–480.
- Rapach, D. E.; J. K. Strauss; and G. Zhou. "International Stock Return Predictability: What Is the Role of the United States?" *Journal of Finance*, 68 (2013), 1633–1662.
- Siriopoulos, C., and A. Fassas. "Implied Volatility Indices: A Review." Working Paper, University of Patras, Greece (2009).
- Stambaugh, R. F. "Predictive Regressions." *Journal of Financial Economics*, 54 (1999), 375–421.
- Ubukata, M., and T. Watanabe. "Market Variance Risk Premiums in Japan as Predictor Variables and Indicators of Risk Aversion." Working Paper, Hitotsubashi University (2011).
- Zhou, G., and Y. Zhu. "A Long-Run Risks Model with Long- and Short-Run Volatilities: Explaining Predictability and Volatility Risk Premium." Working Paper, Washington University in St. Louis (2009).
- Zhou, H. "Variance Risk Premia, Asset Predictability Puzzles, and Macroeconomic Uncertainty." Working Paper, Federal Reserve Board, Washington, DC (2010).