

Expected Stock Returns and Variance Risk Premia*

Tim Bollerslev[†] and Hao Zhou[‡]

First Draft: September 2006

This Version: May 2007

Abstract

We find that the difference between implied and realized variation, or the variance risk premium, is able to explain more than fifteen percent of the ex-post time series variation in quarterly excess returns on the market portfolio over the 1990 to 2005 sample period, with high (low) premia predicting high (low) future returns. The magnitude of the return predictability of the variance risk premium easily dominates that afforded by standard predictor variables like the P/E ratio, the dividend yield, the default spread, and the consumption-wealth ratio (CAY). Moreover, combining the variance risk premium with the P/E ratio results in an R^2 for the quarterly returns of more than twenty-five percent. The results depend crucially on the use of “model-free”, as opposed to standard Black-Scholes, implied variances, and realized variances constructed from high-frequency intraday, as opposed to daily, data. Our findings suggest that temporal variation in both risk-aversion and volatility-risk play an important role in determining stock market returns.

JEL Classification Numbers: G12, G14.

Keywords: Return Predictability, Implied Variance, Realized Variance, Equity Risk Premium, Variance Risk Premium, Time-Varying Risk Aversion.

*Bollerslev's work was supported by a grant from the NSF to the NBER. We would also like to thank John Ammer, Ravi Bansal, Craig Burnside, Pete Kyle, David Lando, Monika Piazzesi, Jonathan Wright, and seminar participants from Federal Reserve Board and Return Predictability Conference at Copenhagen Business School for helpful discussions. The views presented here are solely those of the authors and do not necessarily represent those of the Federal Reserve Board or its staff.

[†]Department of Economics, Duke University, Post Office Box 90097, Durham NC 27708, and NBER, USA, Email bollev@econ.duke.edu, Phone 919-660-1846, Fax 919-684-8974.

[‡]Division of Research and Statistics, Federal Reserve Board, Mail Stop 91, Washington DC 20551 USA, E-mail hao.zhou@frb.gov, Phone 202-452-3360, Fax 202-728-5887.

Expected Stock Returns and Variance Risk Premia

Abstract

We find that the difference between implied and realized variation, or the variance risk premium, is able to explain more than fifteen percent of the ex-post time series variation in quarterly excess returns on the market portfolio over the 1990 to 2005 sample period, with high (low) premia predicting high (low) future returns. The magnitude of the return predictability of the variance risk premium easily dominates that afforded by standard predictor variables like the P/E ratio, the dividend yield, the default spread, and the consumption-wealth ratio (CAY). Moreover, combining the variance risk premium with the P/E ratio results in an R^2 for the quarterly returns of more than twenty-five percent. The results depend crucially on the use of “model-free”, as opposed to standard Black-Scholes, implied variances, and realized variances constructed from high-frequency intraday, as opposed to daily, data. Our findings suggest that temporal variation in both risk-aversion and volatility-risk play an important role in determining stock market returns.

JEL Classification Numbers: G12, G14.

Keywords: Return Predictability, Implied Variance, Realized Variance, Equity Risk Premium, Variance Risk Premium, Time-Varying Risk Aversion.

1 Introduction

Is the return on the stock market predictable? This age-old question arguably still ranks among the most studied and contentious issues in all of economics. Notable recent studies that report evidence in favor of return predictability include Shiller (1984), Keim and Stambaugh (1986), Fama and French (1988a), and Campbell and Shiller (1988b) who find that the ratios of price to earning or price to dividend are useful predictors; Campbell (1991) and Hodrick (1992) who report that the de-trended risk-free rate contains forecasting power; the studies of Fama and Schwert (1977), Campbell (1987) and Fama and French (1989) involving different term structure variables; Lamont (1998) who argues that the aggregate dividend payout ratio forecasts excess returns; and Lettau and Ludvigson (2001) who find the consumption-wealth ratio (CAY) to be a good predictor of future stock market performance; see also the general discussions in support of return predictability in Lewellen (2004) and Cochrane (2006). Even though the predictor variables differ across the above cited studies, most suggest that predictability tends to be the strongest over long multi-year horizons. At the same time, there is also evidence that the degree of predictability appears to have diminished somewhat beginning in the mid-to-late 1990s. In lieu of these findings, we show that the difference between “model-free” implied and realized variances, which we term the variance risk premium, provides remarkable accurate and stable forecasts for the quarterly market return over the recent 1990 to 2005 sample period, with high (low) premia predicting high (low) future returns. Moreover, the accuracy of these forecasts easily dominates that afforded by all of the more traditional predictor variables.

The dual variance concepts underlying our variance risk premium are both fairly new. On the one hand, several recent studies have argued for the use of so-called “model-free realized variances” computed by the summation of high-frequency intraday squared returns.¹ As demonstrated in that literature, these types of measures afford much more accurate *ex-post observations* on the actual return variation than the more traditional sample variances based on daily or coarser frequency return observations (Andersen, Bollerslev, Diebold, and

¹Earlier influential empirical studies employing similar ideas include Schwert (1990) and Hsieh (1991).

Ebens, 2001; Barndorff-Nielsen and Shephard, 2002; Meddahi, 2002). On the other hand, the recently developed so-called “model-free implied variances” provide *ex-ante risk-neutral expectations* of the future return variation. Importantly, and in contrast to the standard option-implied variances based on the Black-Scholes pricing formula or some variant thereof, the “model-free implied variances” are computed from option prices without the use of any particular option-pricing model (Carr and Madan, 1998; Britten-Jones and Neuberger, 2000; Jiang and Tian, 2005).

We show that the difference between these new “model-free” implied and realized variances, which we term the variance risk premium, is able to explain more than fifteen percent of the variation in quarterly stock market returns over the 1990 to 2005 period.² This high degree of predictability clearly dominates that afforded by standard predictor variables like the P/E ratio, the dividend yield, the default spread, the term spread, the risk-free rate, and the consumption-wealth ratio (CAY). Combining the variance risk premium with these other predictor variables, we find that the addition of the P/E ratio results in the greatest overall return predictability of more than twenty-six percent. As explained below, we interpreted this to be consistent with the view that riskiness and risk-aversion are both important for describing the temporal variation in expected stock returns. We further support our empirical results by a number of robustness checks. In this regard we find the use of the “model-free” variance concepts to be crucial, in that the corresponding results based on the traditional Black-Scholes implied variances and realized variances constructed from lower frequency daily data do not support the same strong conclusions.

Why should the difference between implied and realized variances predict stock market returns? Our direct cross-sectional tests offer little support for the idea that the variance risk premium is a systematic risk factor in the usual linear arbitrage pricing (APT) sense. Instead, the variance risk premium may be related (non-linearly) to the coefficient of relative

²Related empirical links between stock market returns and various notions of variance risk have been informally explored by finance professionals. For example, Bondarenko (2004) documents that many equity-oriented hedge funds actively trade variance risk in the highly liquid OTC variance swap market (see, e.g. Gangahar, 2005). Similarly, Beckers and Bouten (2005) report that a market timing strategy based on the ratio of implied to historical volatilities results in doubling the Sharpe ratio relative to that of a constant S&P 500 exposure.

risk aversion within a representative agent setting (Bakshi and Madan, 2006). In a model with time-varying risk aversion, as in, e.g., a model with habit persistence (Campbell and Cochrane, 1999), this therefore provides a direct link for the variance risk premium to explain temporal variation in expected returns. Intuitively, when the variance risk premium is high (low), it generally signals a high (low) degree of risk aversion throughout the economy. Consequently agents tend to simultaneously cut (increase) their consumption and investment expenditures and shift their portfolios from more (less) to less (more) risky assets. This in turn results in a rise (decrease) in expected excess returns and a decrease (increase) in economic growth. In other words, booms and busts follow changes in the market implied risk aversion, as proxied by the difference between the implied and realized variances.

Another theoretical model that could potentially explain our empirical findings is the Bansal and Yaron (2004) long-run risk model, as further extended by Tauchen (2005). In this model volatility risk is explicitly priced despite the fact that innovations in consumption and volatility are uncorrelated. As such, temporal variation in the volatility risk premium will help predict future returns. In particular, while the long-run risk model implies a linear relationship between the price-dividend ratio and the time-varying variance risk premium, our findings indirectly suggest that the financial market based implied and realized variation measures provide much more accurate assessments of the temporal variation in the risk premium than do the *ex post* reported dividends.

Our data sample is admittedly fairly short due to the dual requirements of high-frequency data and liquid options with rich strikes. As such, we can only speculate how well the superior performance of the variance risk premium relative to the other macroeconomic-finance predictor variables will hold up in the future and over longer time periods. Nonetheless, one aspect of our results seems particularly robust to the many different settings and appears to hold up internationally as well. Namely, the difference between the implied and realized variances, or the variance risk premium, always works better in predicting the returns than do any one of the two variation measures in isolation. This result in turn may help guide the formulation of more structural tests for specific economic models and hypothesis—whether it is time-varying risk aversion (Campbell and Cochrane, 1999) or time-varying volatility risk

(Bansal and Yaron, 2004) that primarily drives the apparent stock return predictability.³

The plan for the rest of the paper is as follows. Section 2 discusses the concepts of “model-free” implied and realized variances used in defining the variance risk premium, along with related practical data issues. Section 3 presents our main results on return predictability and corresponding robustness checks. Section 4 concludes.

2 Variance Risk Premium

2.1 Formal Definitions

In parallel to the standard definition of a forward premium, the variance risk premium is defined as the difference between a current return variation measure and the market’s expectation of some future return variation. Formally, let $C_t(T, K)$ denote the price of a European call option maturing at time T with strike price K , and $B(t, T)$ denote the price of a time t zero-coupon bond maturing at time T . As shown by Carr and Madan (1998), Demeterfi, Derman, Kamal, and Zou (1999) and Britten-Jones and Neuberger (2000), the implied variance or the market’s risk-neutral expectation of the total return variation between time t and $t + 1$ conditional on time t information, say IV_t , may then be expressed in a “model-free” fashion as the following portfolio of European calls of the underlying futures price,

$$\begin{aligned}
 IV_t &\equiv \sum_{i=1}^m \left[\frac{C_t\left(t+1, \frac{K_i}{B(t,t+1)}\right) - C_t\left(t, \frac{K_i}{B(t,t)}\right)}{K_i^2} - \frac{C_t\left(t+1, \frac{K_{i-1}}{B(t,t+1)}\right) - C_t\left(t, \frac{K_{i-1}}{B(t,t)}\right)}{K_{i-1}^2} \right] \Delta K \\
 &\longrightarrow 2 \int_{K_0}^{K_m} \frac{C_t\left(t+1, \frac{K}{B(t,t+1)}\right) - C_t\left(t, \frac{K}{B(t,t)}\right)}{K^2} dK, \\
 &\longrightarrow E_t^Q[\text{Return Variation}(t, t+1)],
 \end{aligned} \tag{1}$$

where $\Delta K = (K_m - K_0)/m$ and $K_i = K_0 + i\Delta K$, and $m \rightarrow \infty$ for the second equation to hold, while $K_m \rightarrow \infty$ and $K_0 \rightarrow 0$ for the third equation to hold; i.e., an ever increasing

³Of course, from a purely empirical perspective it is difficult to distinguish between the hypothesis of a volatility risk premium as a predictor for the returns and the notion of time-varying risk aversion.

number of calls with strikes spanning zero to infinity.⁴ Importantly, however, as shown in Jiang and Tian (2005) and Bollerslev, Gibson, and Zhou (2006), in practice the “model-free” implied variance generally provides an accurate approximation to the true (unobserved) risk-neutral expectation of the return variation with relatively few different strikes (m), and, in particular, a much better approximation than the one afforded by inversion of the standard Black-Scholes formula.

In order to define the actual variance, let p_t denote the logarithmic price of the asset. The realized variation over the discrete $t - 1$ to t time interval is then naturally measured in a “model-free” fashion by

$$RV_t \equiv \sum_{j=1}^n \left[p_{t-1+\frac{j}{n}} - p_{t-1+\frac{j-1}{n}(\Delta)} \right]^2 \longrightarrow \text{Return Variation}(t-1, t) \quad (2)$$

where the convergence relies on $n \rightarrow \infty$; i.e., an increasing number of within period price observations. In practice, market microstructure frictions invariably limit the highest feasible sampling frequency. For liquid assets, a five-minute frequency usually strikes a reasonable balance between the desire for as finely sampled returns as possible and the “contaminating” influences of price discreteness, the bid-ask spread and other frictions. As demonstrated in the literature (Andersen, Bollerslev, Diebold, and Ebens, 2001; Barndorff-Nielsen and Shephard, 2002; Meddahi, 2002), this type of “model-free” variance measure based on high-frequency intraday data affords much more accurate ex-post observations of the actual return variation than the more traditional sample variances based on daily or coarser frequency returns.

The variance risk premium underlying our empirical investigations is simply defined as the difference between the ex-ante risk neutral expectation of the future return variation over the $[t, t + 1]$ time interval and the ex-post realized return variation over the $[t - 1, t]$

⁴The expression in equation (1) is a special case of the “model-free” forward implied variance from T_1 to T_2 conditional on time t information for general jump-diffusion processes (Bondarenko, 2004; Jiang and Tian, 2005; Carr and Wu, 2005),

$$IV_t \equiv 2 \int_0^\infty \frac{C_t \left(T_2, \frac{K}{B(t, T_2)} \right) - C_t \left(T_1, \frac{K}{B(t, T_1)} \right)}{K^2} dK = E_t^Q [\text{Return Variation} (T_1, T_2)].$$

Letting $T_1 = t$ and $T_2 = t + 1$, results in the one-period, or spot, implied variance defined in the text.

time interval,

$$VRP_t \equiv IV_t - RV_t. \tag{3}$$

Note that IV_t and RV_t , and therefore VRP_t , are both directly observable at time t . This is, of course, important from a forecasting perspective. However, this definition of the variance risk premium differs slightly from the one used by some other recent studies, involving the difference between the variation measures over the identical time intervals, or in terms of the present notation the ex-post variance spread, $IV_t - RV_{t+1}$.⁵ In particular, Bollerslev, Gibson, and Zhou (2006) find that the temporal variation in $IV_t - E_t(RV_{t+1})$ for the aggregate market portfolio, as implied by a standard Heston (1993) one-factor stochastic volatility model, may in part be explained by a set of macro-finance variables, including lagged realized volatility, the P/E ratio, and the growth rate in industrial production. Similarly, Todorov (2006) has recently explored the dynamics of $IV_t - E_t(RV_{t+1})$ within the context of a very general continuous time specification allowing for separate jump and diffusive risk premiums. Meanwhile, Carr and Wu (2005) find that the common Fama and French (1993) risk factors are able to explain only a small fraction of the ex-post variation in the variance risk premia for a set of individual stocks. The difference between implied and realized variances has also previously been associated empirically with a measure of the aggregate market risk aversion by Rosenberg and Engle (2002). Also, Bakshi and Madan (2006) have formally shown that the volatility spread may be expressed as a nonlinear function of the aggregate degree of risk aversion within a representative agent setting.

2.2 Data Description

Our empirical analysis is based on data for the aggregate S&P500 composite index from January 1990 through January 2005. We rely on monthly data for the VIX index for quantifying the risk-neutral implied variance measure. The VIX index is based on the highly

⁵More generally, a maturity-specific variance risk premium, or variance spread, may be defined as the difference between the risk-neutral and objective expectations of the quadratic variation over the $[t, T]$ time interval, $E_t^Q(RV_{t,T}) - E_t(RV_{t,T})$, where $RV_{t,T}$ denotes the time t to T variation as defined by the limit in equation (2).

liquid S&P500 index options along with the “model-free” approach discussed above explicitly tailored to replicate the risk-neutral variance of a fixed 30-day maturity. The data is obtained directly from the Chicago Board of Options Exchange (CBOE).⁶ Our monthly realized variance measure is based on the summation of the five-minute squared returns for the S&P500 composite index within the month. The high-frequency data for the S&P500 index is provided by the Institute of Financial Markets. For a typical month with 22 trading days there are a total of $n = 22 \times 78 = 1,716$ “five-minute” returns, corresponding to the 78 five-minute subintervals covering the normal trading hours from 9:30am to 4:00pm along with the close-to-open return.

Even though the monthly realized volatilities and the VIX index are both available on an overlapping daily basis, most of the other variables that we consider in our return predictions are only available monthly or quarterly. Thus, to avoid the dual problems associated with overlapping data samples and missing observations, we will primarily focus on quarterly time series and forecast horizons, appropriately using the values for the last day or the last month in each quarter as the “quarterly” observations.⁷

To illustrate, Figure 1 plots the time series of implied variances, realized variances, and their differences on a quarterly basis. As immediately evident from the figure, both of the variance measures are somewhat higher during the 1997 to 2002 part of the sample. The few distinct spikes in the measures generally also tend to coincide. Consistent with the earlier empirical evidence cited above, the spread between the implied and realized variances is always positive, indicative of a systematic variance risk premium.

In addition to the variance risk premium, we also consider a set of other more traditional predictor variables. In particular, we obtain monthly P/E ratios and dividend yields for the S&P500 directly from Standard & Poor’s. Data on the 3-month T-bill, the default spread (between Moody’s BAA and AAA corporate bond spreads), the term spread (between the 10-year T-bond and the 3-month T-bill yields), and the stochastically de-trended risk-free

⁶The CBOE replaced the old VIX index, based on S&P100 options and Black-Scholes implied volatilities, with the new VIX index, based on S&P500 options and “model-free” implied volatilities, in September 2003. Historical data on both indexes are available from the official CBOE website.

⁷For recent discussions of some of the difficulties associated with the use of overlapping data see, e.g., Valkanov (2003) and Boudoukh, Richardson, and Whitelaw (2006).

rate (the 1-month T-bill rate minus its backward 12-month moving averages) are taken from the public website of the Federal Reserve Bank of St. Louis. The consumption-wealth ratio (CAY), as defined in Lettau and Ludvigson (2001) and updated until the first quarter of 2005, is downloaded from Lettau and Ludvigson's website. Similar sets of predictor variables for other forecast comparison purposes and sample periods have been used by Lamont (1998), Lettau and Ludvigson (2001), Ang and Bekaert (2006), among many others.

Basic summary statistics for each of the predictor variables are given in Table 1. The mean excess return on the S&P500 over the sample equals 5.97 percent annually. The sample means for the implied and realized variances are 36.30 and 14.90, respectively, corresponding to a variance risk premium of 21.40 (in percentages squared). The numbers for the traditional forecasting variables are all directly in line with those reported in previous studies. In particular, all of the variables are highly persistent with first order autocorrelations ranging from 0.81 to 0.96. In contrast, the serial correlation in the implied and realized quarterly variances are both much less, and the first order autocorrelation for their difference only equals 0.31. As such, this alleviates one of the common concerns related to the use of highly persistent predictor variables and the possibility of spurious or unbalanced regressions.⁸ Further anticipating the forecasting results discussed next, the traditional predictor variables all correlate fairly weakly with the contemporaneous excess returns with sample correlations ranging from -0.12 to +0.04, while the sample correlations between the returns and the different variance measures are much higher (in an absolute sense) ranging from -0.57 to -0.28.

3 Forecasting Stock Market Returns

All of our forecasts are based on simple linear regressions of the quarterly S&P500 excess returns on the different sets of lagged predictor variables. We focus our comparisons of the accuracy of the different forecasts on the standard adjusted R^2 's. To allow for serial

⁸Inference issues related to the use of highly persistent predictor variables have been studied extensively in the literature, see, e.g., Stambaugh (1999), Ferson, Sarkissian, and Simin (2003), and Campbell and Yogo (2006) and the references therein. Some authors have gone as far as to attribute the apparent predictability to the use of strongly serially correlated predictors, e.g., Goyal and Welch (2003, 2005).

correlation and heteroskedasticity in the residuals from the estimated regressions we rely on Newey and West (1987) robust t -statistics based on four lags.

3.1 Main Results

Table 2 details our main empirical findings. The regression coefficient for the variance risk premium is estimated to be positive. At an intuitive level, this is consistent with the idea that when the market anticipates high (low) volatility going forward, there is a discount (premium) built into prices, in turn resulting in high (low) future returns. Moreover, this simple difference between the implied and realized variance is able to explain as much as 15.14 percent of the quarterly variation in the excess returns. This strong degree of predictability is all the more impressive when compared to the results obtained by regressing the returns on the implied or realized variance in isolation. The corresponding R^2 's, reported in the next two columns, indicate that these account for just 6.32 percent and a negative adjusted -1.05 percent of the variation, respectively.⁹

Comparing the regression results for the variance risk premium to the ones for the traditional predictor variables further underscores the significance of the fifteen percent R -square. In line with the earlier findings of Fama and French (1988a) and Campbell and Shiller (1988a), the logarithmic P/E and P/D ratios only explain 6.22 and 2.76 percent of the variation in the quarterly returns, respectively, while the term spread, the default spread, and the relative risk-free rate all have close to zero or negative R^2 's. Further, the consumption-wealth ratio of Lettau and Ludvigson (2001) accounts for a modest 4.83 percent of the quarterly return variation, although it still dominates most of the other financial or macroeconomic variables. Looking at the robust t -statistics, the variance spread also appears highly statistically significant (3.94), while with the possible exception of IV_t (-2.81), CAY (2.29) and P/E (-2.08), each of the other predictor variables are at best marginally significant at conventional levels.

A number of the studies cited above have, of course, argued that the degree of pre-

⁹This result is also consistent with the findings in Guo and Whitelaw (2006) that implied volatility is more informative than realized volatility.

dictability afforded by the classic P/E and P/D ratios appear to be the strongest over longer multi-year horizons. Our limited data sample (the VIX index is only available from 1990 and onwards) prevents us from looking at the predictability of the variance risk premium over such horizons. At the same time, several empirical studies also suggest that the evidence in favor of predictability seems to have diminished somewhat during the stock market boom of the mid- to late-1990s. The results for the variance risk premium reported here, thus directly complement these findings by providing strong evidence for quarterly return predictability over the recent 1990 to 2005 sample period.

To further illustrate the regression results, we show in Figure 2 the lagged (standardized) time $t - 1$ variance premium, P/E ratio, and CAY, along with the time t quarterly S&P500 returns. Looking at the lower two panels, the P/E ratio and CAY both seem to capture the longer cycles in the quarterly excess returns, with CAY being slightly more responsive to short-run fluctuations as well. Meanwhile, the difference between the implied and the realized variances depicted in the top panel not only tracks the longer-run cycles in the excess returns reasonably well, but it also seems to lead the quarter-to-quarter return variation. This latter pattern may be especially pronounced for the marked spikes associated with the 1991 recession, the 1998 Russian debt and LTCM crises, and the 2002 corporate accounting scandals.

In addition to the simple regressions discussed above, the right panel of Table 2 reports the results from a representative set of multiple regressions. Combining the variance premium with the P/E ratio results in an impressive R^2 of 26.37 percent. Such a high R^2 clearly surpasses the forecasting power for quarterly returns previously reported in the literature with any other forecasting variables. Interestingly, this R^2 is also much higher than the sum of the R^2 's from the two individual forecasting regressions for VRP_t and P/E. Intuitively, if expected stock returns are driven by temporal variation in both risk and risk-aversion, a combination of the P/E ratio and the variance risk premium may effectively account for both. Alternatively, the P/E ratio and the variance risk premium may jointly capture both the long-run and short risks embedded in the stock market returns. A similar pattern, although to a lesser degree, obtains when including the CAY variable along with the variance

premium, which produces an R^2 of 20.86 percent. On the other hand, combining the P/E ratio and CAY, the adjusted R^2 is merely 5.54 percent, and the robust t -statistics turn out to be insignificant for both. Also, even though the term spread and the relative short rate are insignificant by themselves, both variables contribute marginally to a joint predictive regression with the variance premium and the P/E ratio, resulting in the highest overall adjusted R^2 of 27.67 percent.¹⁰ Lastly, it is worth noting that the estimated coefficients for the variance risk premium remain remarkably stable, ranging from 0.87 to 1.11 with corresponding t -statistics between 3.85 and 5.27, regardless of the other variables included in the forecast regressions.

3.2 Robustness Check

These striking empirical findings naturally raise questions about the robustness and the validity of the results in a wider sense. We next discuss a series of additional regressions and sensitivity checks designed to address these concerns.

3.2.1 Alternative Variance Measures

The “model-free” implied and realized variances are both relatively new concepts. As a first robustness check, we replace the “model-free” variance measures with the standard Black-Scholes implied variance, IV_t^* , and the realized variance based on low-frequency daily returns, RV_t^* . The regression results for these “old” variance measures are reported in Table 3. The variance risk premium defined by the difference between the Black-Scholes implied variance and the daily return based realized variance still dominates the results based on each of the variance measures in isolation (R^2 of 3.62 percent versus 3.05 and 0.68 percent, respectively). Meanwhile, replacing the Black-Scholes implied variance with the “model-free” implied variance increases the predictability to 8.08 percent. Similarly, substituting the realized variance based on daily data with the corresponding high-frequency based measure increases the R^2 to 7.30 percent.

¹⁰This latter finding is broadly consistent with the earlier results in Fama and French (1989), Campbell (1991), and Hodrick (1992) related to the usefulness of term structure information for predicting stock returns.

With the exception of P/E and CAY, the R^2 for $IV_t^* - RV_t^*$ is still higher than those for all of the other macro-finance variables included in Table 2. However, the R^2 for the “model-free” variance premium in Table 2 (15.14 percent) is orders of magnitude greater than the one based on the “old” variance measures (3.62 percent). Combining the “old” variance premium and the P/E ratio in the same regression results in an R^2 of 14.35 percent, compared to 26.37 percent for the “model-free” measures. Similarly, mixing the “old” and the “new” variance measures in the extended regressions results in lower R^2 's of 20-22 percent and 14-16 percent for $IV_t^* - RV_t$ and $IV_t - RV_t^*$, respectively.

All-in-all, comparing the results across Tables 2 and 3, clearly highlight the superior predictive performance afforded by the “model-free” measures and corresponding variance risk premium.¹¹

3.2.2 Monthly Return Predictability

As previously noted, for most predictor variables the evidence in favor of return predictability tend to be the strongest over longer multi-year horizons. The limited availability of the “model-free” variance measures prevents us from meaningfully looking beyond a quarter. Instead, as a further robustness check we report in Table 4 the results from a set of non-overlapping monthly return regressions. Not surprisingly, the degree of predictability afforded by all of the regressions, as measured by the R -square, is in the low single digit. Interestingly, however, with the exception of the P/E ratio, the “model-free” variance risk premium dominates all of the other macro-finance variables, and also has the most significant t -statistic. Moreover, the R^2 for the difference between the two variance measures (1.24%) clearly outperforms the implied variance (0.58%) and the realized variance (-0.42%) in isolation. Combining the variance risk premium with the P/E ratio again results in an R^2 (3.98%) in excess of the sum of the two R^2 's from the individual regressions (1.24% and 2.10%), indirectly pointing to the importance of capturing temporal variation in both risk and risk-aversion. Of course, assuming that the degree of risk aversion changes only slowly,

¹¹Of course, it is possible that more sophisticated parametric volatility models may result in even better predictions than the ones afforded by the “model-free” measures, but care need to be taken to avoid any “look-ahead” bias and deal with parameter estimation error uncertainty.

it is hardly surprising that predictability at the monthly horizon is significantly lower than the previously reported quarterly R^2 's. The monthly horizon and corresponding unrestricted regressions may simply be too short to capture the long-run risk component embedded in the market returns.

3.2.3 Volatility Risk Premium

The regression results discussed above have all been based on implied and realized variances, and their difference. This mirrors the common use of variance denominated contracts in the over-the-counter swap market (Demeterfi, Derman, Kamal, and Zou, 1999).¹² From a theoretical perspective, however, the volatility, or any other nonlinear monotone transformation of the variance, may serve as an alternative, and depending on the situation more appropriate measure of risk (Merton, 1980). To this end, we report in Table 5 the results obtained by replacing the “model-free” variance measures with their volatility, or standard deviation, counterparts. Interestingly, including the volatility difference, $\sqrt{IV_t} - \sqrt{RV_t}$, in isolation produces an even higher predictive R^2 of 18.50 percent, compared to the 15.14 percent for the variance premium. Again, including the implied and realized volatility by themselves results in adjusted R^2 's close to zero. Also, the coefficient estimates for the volatility difference are very stable and between 4.07 to 4.39, with t -statistics around 4.5, in both the simple and multiple regressions. Nonetheless, the predictive performance in the multiple regressions obtained by including the P/E ratio, CAY, and the term structure variables all fall short of the comparable results for the variance premium in Table 2.

3.2.4 Alternative Predictor Variables

The different variables included in the previous tables arguably comprise some of the most commonly advocated predictors. However, several other variables have, of course, been proposed in the literature. While it is literally impossible to investigate all of these suggestions, we report in Table 6 the results for a few additional predictor variables and regressions.¹³

¹²The CBOE futures on the VIX, which began trading in 2003, are also in variance form, as is the realized variance futures on the S&P500, which have been trading since 2004; see also Mixon (2007).

¹³We also experimented with the bond factor of Cochrane and Piazzesi (2005), but found that it had little predictability with the 1990 to 2005 quarterly returns analyzed here.

First, a number of studies have argued that returns are subject to mean reversion, or negative serial correlation, although again mostly so over longer multi-year horizons; see, e.g., Fama and French (1988b) and Lo and MacKinlay (1988). The estimates in the first column show that the lagged quarterly excess return has no forecasting power over the present sample period.

Second, Lamont (1998) finds that the corporate payout ratio is useful for predicting quarterly returns. Consistent with the findings in Lettau and Ludvigson (2001), the estimates in the second column indicate that this predictability has completely disappeared over the more recent 1990 to 2005 sample period.

Third, as previously noted, Bollerslev, Gibson, and Zhou (2006) have recently presented forward looking estimates of the volatility risk premium, λ_t , within the context of the Heston (1993) one-factor stochastic volatility model, where λ_t is allowed to be time-varying and dependent on important macro-finance variables. The third column reports the regression results using this model-specific parametrically estimated premium. Although the R^2 of 4.47 percent is on par with that of CAY in the simple regression, the corresponding robust t -statistic of 1.66 is at best marginally significant.¹⁴ Moreover, including λ_t along with the other explanatory variables, completely overturns this significance.

Lastly, we report the results obtained by including the so-called realized quarticity measure

$$RQ_t \equiv \sum_{j=1}^n \left[p_{t-1+\frac{j}{n}} - p_{t-1+\frac{j-1}{n}(\Delta)} \right]^4 .$$

This measure (or a proportional transformation thereof) may be seen as an estimate for the variance of the return variation (Barndorff-Nielsen and Shephard, 2002), or the volatility-of-volatility. Interestingly, not only is the realized quarticity significant by itself, but it remains so when included together with the variance risk premium and the other explanatory variables as indicated in the last column. Indeed, the quarticity adds an additional 2.5 percent to the R^2 when compared to the best result in Table 2, suggesting that further (marginal) forecast improvements might even be possible by considering the variance-of-

¹⁴Since the estimation of λ_t is based on the full sample, these results are also be subject to a “look ahead” bias.

variance risk premium.

Taken as a whole, the results in Table 6 again underscore the robustness and superior predictability of the new “model-free” variance risk premium measure, whether combined with the P/E ratio and some of the other predictor variables, or when used in isolation.

3.2.5 International Evidence

The length of our sample is admittedly fairly short. Reliable high quality options and intraday data for the S&P500 index only became available in the late 1980s, and there isn’t much we can do to extend the sample prior to that time. Instead, we here present some limited international evidence based on German data. In particular, we construct quarterly German DAX index returns, corresponding model-free implied variance measures (VDAX-NEW), realized variance measures based on daily DAX returns, and German three-month Treasury bill rates, for the sample period extending from 1993Q3 to 2007Q1.¹⁵

The resulting quarterly forecasting regression show that the DAX variance premium explains 4.16 percent of the variation in the quarterly German returns, with an estimated slope coefficient of 0.87 and a t-statistic of 2.54. Albeit not as significant, these results are broadly consistent with the ones for the S&P500 based on realized variances constructed from daily returns reported in the third column in Table 3. More importantly, however, the same basic finding that the variance difference is a much better predictor than the two variation measures in isolation remains intact. Including the implied or the realized variance series one at a time in quarterly DAX regressions results in adjusted R^2 's of -1.42 and -1.83, respectively. Of course, the liquidity of the DAX options are inferior to those for the S&P500, and the German results necessarily also reflect a number of other country specific differences. Nonetheless, we are encouraged by the general coherence between the results across the two countries.

¹⁵The data were obtained from several sources including Deutsche Börse, Bloomberg, Yahoo Finance, and the Federal Reserve Board. Further details are available upon request.

3.3 Economic Interpretation

The striking empirical findings reported above naturally begs the question of why the implied-realized variance difference has the power to forecast returns? The most likely answer is that the premium proxies for time-varying risk aversion. There already exists a large literature concerned with ways in which to extract the market implied risk aversion from options data (see, e.g., Jackwerth, 2000; Ait-Sahalia and Lo, 2000; Rosenberg and Engle, 2002; Balyeat, 2002; Brandt and Wang, 2003; Bliss and Panigirtzoglou, 2004; Gordon and St-Amour, 2004; Wu, 2005; Garcia, Lewis, Pastorello, and Renault, 2006). Most of these studies rely on specific parametric models (like Black-Scholes) for inferring the risk-neutral expectations, while generally also resorting to fairly complex and numerical intensive procedures for estimating the risk-neutral distribution. In contrast, the risk-neutral expected variance used here is based on the idea of replication using a portfolio of European call options sufficiently rich along the strike dimension. The wedge between this “model-free” market-based expectation and the actual return variation thus serves as a measure of the difference between the risk-neutral and objective distributions, or a measure of the market implied risk aversion. The fact that the P/E ratio – as a valuation measure involving both risk and risk aversion – adds more than one-to-one to the degree of predictability relative to the spread alone, also indirectly suggests that the variance premium isn’t simply measuring market risk.¹⁶

To further appreciate the economics behind the apparent connection between the variance risk premium and returns, Figure 3 plots the lagged premium together with the quarterly growth rate in GDP. As seen from the figure, there is a tendency for the variance risk premium to rise in the quarter before a decline in GDP, while it typically narrows ahead of an increase in GDP. Indeed, the sample correlation between the two series equals -0.16. Thus, not only is the difference between the implied and realized variances positively correlated with expected future returns, it also covaries negatively with the future growth rate in GDP.¹⁷ Intuitively, when the variance risk premium is high (low), it generally signals a high (low) degree of risk

¹⁶Of course, the P/E ratio may also reflect expected earning growth, the term structure, the payout ratio, and potentially a bubble component; which may or may not be fully captured by just risk and risk aversion.

¹⁷This short-run negative correlation between the variance risk premium and GDP growth rate mirrors the longer-run counter cyclical behavior of CAY noted by Lettau and Ludvigson (2001).

aversion throughout the economy. Consequently agents tend to simultaneously cut (increase) their consumption and investment expenditures and shift their portfolios from more (less) to less (more) risky assets. This in turn results in a rise (decrease) in expected excess returns and a decrease (increase) in economic growth. In other words, booms and busts are directly associated with changes in the market implied risk aversion, as proxied by the difference between the implied and realized variances.¹⁸

To further underscore that the variance risk premium is best seen as a proxy for risk aversion as opposed to systematic risk, we complement the time series based evidence with standard cross-sectional asset pricing tests, examining whether the factor loadings for a set of portfolios with different risk characteristics are appropriately dispersed and whether the estimated loadings are actually statistically significant (Daniel and Titman, 2005; Lewellen and Nagel, 2006; Lewellen, Nagel, and Shanken, 2006; Burnside, 2006). More specifically, we begin by running the time-series regressions

$$R_{it}^e = \alpha_i + \beta_i \times \text{Factor}_t + u_{it} ,$$

where R_{it}^e denotes the monthly excess return on the twenty-five Fama and French (1992) size and book-to-market sorted portfolios, and Factor_t refers to the market excess return and the variance risk premium, respectively.¹⁹ The resulting β_i estimates for the market returns, reported in the top panel of Table 7, are reasonably spread out (ranging from 0.68 to 1.47), while the estimated loadings for the variance risk premium are concentrated in a much narrower range (from -1.14 to -0.81). Moreover, the t -statistics for the market factor are all highly statistically significant, while only about half of the loadings for the variance premium are significant at the 5% level. This narrow spread in the estimated loadings coupled with the lack of statistical significance indirectly suggest that the variance risk premium is not a priced risk factor (see, e.g., Burnside, 2006). The results from the corresponding cross-

¹⁸The model of Campbell and Cochrane (1999) also generates such a counter cyclical risk aversion for the marginal investor. In a related context Campbell and Diebold (2006) document that expected business conditions, as measured by the Livingston survey data, are indeed negatively correlated with expected returns.

¹⁹Note, both the market excess return and the variance risk premium are directly traded and have zero mean under the risk-neutral expectation.

sectional regressions,

$$\bar{R}_i^e = \hat{\beta}_i \times \lambda + u_i ,$$

where \bar{R}_i^e refers to the sample mean returns and $\hat{\beta}_i$ denotes the estimated loadings, further supports this conjecture. The risk premium estimate for the market risk factor is positive and statistically significant, while the estimated λ for the variance factor is negative and insignificant. Thus, even though the relatively short time span (1990-2005) invariably limits the capacity of these cross-sectional tests to effectively discriminate among competing risks, the results are entirely consistent with the idea that the strong time series predictability associated with the variance risk premium is primarily due to temporal variation in risk aversion, as opposed to time-varying systematic risk.²⁰

One structural model that could potentially explain our empirical findings is the Bansal and Yaron (2004) long-run risk model and its extension by Tauchen (2005). Under the assumption of recursive utility the long-run risk model implies that volatility carries a risk premium, despite the fact that innovations in consumption and volatility are uncorrelated. As such, fluctuations in stock returns are predictable by temporal variation in the volatility risk premium. In particular, the long-run risk model posits that the logarithmic price-dividend ratio for the aggregate market is a linear function of long-term growth risk, say x_t , and time-varying volatility risk, say σ_t^2 ,

$$\log(P_t/D_t) = A_0 + A_1 x_t + A_2 \sigma_t^2 .$$

In light of this relationship, our empirical findings may therefore be interpreted as indirectly supporting the idea that the market-based variance risk premium provides a more accurate measure of the temporal variation in the return premium than do the reported dividends. Moreover, combining the P/E ratio and the variance risk premium may effectively capture both long-run and short-run predictable components in the returns. Empirically, of course, it is difficult to distinguish between the hypothesis of a volatility risk premium and the notion of time-varying risk aversion. Nonetheless, it is our hope that the results reported

²⁰Recent studies by Ang, Hodrick, Xing, and Zhang (2006) and Adrian and Rosenberg (2006) have argued that other measures of “volatility” might be priced risk factors.

here will prove useful in the design of new tests and empirical procedures to more forcefully distinguish between these hypotheses.

4 Conclusion

We provide strong statistical evidence that stock market returns are predictable by the difference between implied and realized variances, or the variance risk premium. The economic magnitude of the documented predictability is large, with the variance risk premium accounting for more than fifteen percent of the quarterly return variation. Combining the variance risk premium with a set of other traditional explanatory variables explains close to thirty percent of the ex-post return variation. These results are robust across different specifications and predictor variables and also appear to hold internationally.

The 1990-2005 sample period underlying most of our empirical investigations experienced quite large return fluctuations. These fluctuations are difficult to justify by variations in economic fundamentals and risk alone. Instead, the variance risk premium may be seen as a “model-free” measure for the aggregate degree of risk aversion in the economy and/or a proxy for volatility risk. As such, our empirical results suggest that both time-varying risk-aversion (Campbell and Cochrane, 1999) *and* time-varying volatility risk (Bansal and Yaron, 2004) may all be important for understanding temporal variations in expected stock market returns. It would be interesting, but beyond the scope of the present paper, to more structurally discriminate between these competing explanations.

References

- Adrian, Tobias and Joshua Rosenberg (2006), “Stock Returns and Volatility: Pricing the Long-Run and Short-Run Components of Market Risk,” Working Paper, Federal Reserve Bank of New York.
- Aït-Sahalia, Yacine and Andrew W. Lo (2000), “Nonparametric Risk Management and Implied Risk Aversion,” *Journal of Econometrics*, vol. 94, 9–51.
- Andersen, Torben G., Tim Bollerslev, Francis X. Diebold, and Heiko Ebens (2001), “The Distribution of Realized Stock Return Volatility,” *Journal of Financial Economics*, vol. 61, 43–76.
- Ang, Andrew and Geert Bekaert (2006), “Stock Return Predictability: Is it There?” *Review of Financial Studies*, Forthcoming.
- Ang, Andrew, Robert Hodrick, Yuhang Xing, and Xiaoyan Zhang (2006), “The Cross-Section of Volatility and Expected Returns,” *Journal of Finance*, vol. 61, 259–299.
- Bakshi, Gurdip and Dilip Madan (2006), “A Theory of Volatility Spread,” Working Paper, Robert H. Smith School of Business, University of Maryland.
- Balyeat, R. Brian (2002), “Economic Significance of Risk Premiums in the S&P500 Option Market,” *Journal of Futures Markets*, vol. 22, 1147–1178.
- Bansal, Ravi and Amir Yaron (2004), “Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles,” *Journal of Finance*, pages 1481–1509.
- Barndorff-Nielsen, Ole and Neil Shephard (2002), “Econometric Analysis of Realised Volatility and Its Use in Estimating Stochastic Volatility Models,” *Journal of Royal Statistical Society*, vol. Series B 64.
- Beckers, Stan and Tom Bouten (2005), “Implied Volatility and Market Timing,” *Finance Letters*, vol. 3, 12–16.
- Bliss, Robert R. and Nikolaos Panigirtzoglou (2004), “Option-Implied Risk Aversion Estimates,” *Journal of Finance*, vol. 59, 407 – 446.
- Bollerslev, Tim, Mike Gibson, and Hao Zhou (2006), “Dynamic Estimation of Volatility Risk Premia and Investor Risk Aversion from Option-Implied and Realized Volatilities,” Working Paper, Federal Reserve Board.
- Bondarenko, Oleg (2004), “Market Price of Variance Risk and Performance of Hedge Funds,” *Working Paper*, Department of Finance, University of Illinois.

- Boudoukh, Jacob, Matthew Richardson, and Robert F. Whitelaw (2006), “The Myth of Long-Horizon Predictability,” NBER Working Paper No.11841.
- Brandt, Michael W. and Kevin Q. Wang (2003), “Time-varying Risk Aversion and Expected Inflation,” *Journal of Monetary Economics*, vol. 50, 1457–1498.
- Britten-Jones, Mark and Anthony Neuberger (2000), “Option Prices, Implied Price Processes, and Stochastic Volatility,” *Journal of Finance*, vol. 55, 839–866.
- Burnside, Craig (2006), “Robust Estimation and Inference for Stochastic Discount Factor Models: Avoiding Nocuous Normalizations and Spurious Risk Factors,” *Working Paper*, Department of Economics, Duke University.
- Campbell, John Y. (1987), “Stock Returns and the Term Structure,” *Journal of Financial Economics*, vol. 18, 373–399.
- Campbell, John Y. (1991), “A Variance Decomposition for Stock Returns,” *Economic Journal*, vol. 101, 157–179.
- Campbell, John Y. and John H. Cochrane (1999), “By Force of Habit: A Consumption Based Explanation of Aggregate Stock Market Behavior,” *Journal of Political Economy*, vol. 107, 205–251.
- Campbell, John Y. and Robert J. Shiller (1988a), “The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors,” *Review of Financial Studies*, vol. 1, 195–228.
- Campbell, John Y. and Robert J. Shiller (1988b), “Stock Prices, Earnings, and Expected Dividends,” *Journal of Finance*, vol. 43, 661–676.
- Campbell, John Y. and Motohiro Yogo (2006), “Efficient Tests of Stock Return Predictability,” *Journal of Financial Economics*, vol. 81, 27–60.
- Campbell, Sean D. and Francis X. Diebold (2006), “Stock Returns and Expected Business Conditions: Half a Century of Direct Evidence,” *Working Paper*, University of Pennsylvania.
- Carr, Peter and Dilip Madan (1998), *Towards a Theory of Volatility Trading*, chap. 29, pages 417–427, Risk Books.
- Carr, Peter and Liuren Wu (2005), “Variance Risk Premia,” *Working Paper*, NYU and Baruch College at CUNY.

- Cochrane, John H. (2006), “The Dog That Did Not Bark: A Defense of Return Predictability,” NBER Working Paper No.12026.
- Cochrane, John H. and Monika Piazzesi (2005), “Bond Risk Premia,” *American Economic Review*, vol. 95, 138–160.
- Daniel, Kent and Sheridan Titman (2005), “Testing Factor-Model Explanations of Market Anomalies,” *Working Paper*, Northwestern University.
- Demeterfi, Kresimir, Emanuel Derman, Michael Kamal, and Joseph Zou (1999), “A Guide to Volatility and Variance Swaps,” *Journal of Derivatives*, vol. 6, 9–32.
- Fama, Eugene and Kenneth French (1992), “The Cross-Section of Expected Stock Returns,” *Journal of Finance*, vol. 47, 427–465.
- Fama, Eugene and Kenneth French (1993), “Common Risk Factors in the Returns on Stocks and Bonds,” *Journal of Financial Economics*, vol. 33, 3–56.
- Fama, Eugene F. and Kenneth R. French (1988a), “Dividend Yields and Expected Stock Returns,” *Journal of Financial Economics*, vol. 22, 3–25.
- Fama, Eugene F. and Kenneth R. French (1988b), “Permanent and Temporary Components of Stock Prices,” *Journal of Political Economy*, vol. 96, 246–273.
- Fama, Eugene F. and Kenneth R. French (1989), “Business Conditions and Expected Returns on Stocks and Bonds,” *Journal of Financial Economics*, vol. 25, 23–49.
- Fama, Eugene F. and G. William Schwert (1977), “Asset Returns and Inflation,” *Journal of Financial Economics*, vol. 5, 115–146.
- Person, Wayne E., Sergei Sarkissian, and Timothy T. Simin (2003), “Spurious Regressions in Financial Economics?” *Journal of Finance*, pages 1393–1414.
- Gangahar, Anuj (2005), “Volatility Becomes an Asset Class,” *Financial Times*, May 23rd Issue.
- Garcia, René, Marc-André Lewis, Sergio Pastorello, and Éric Renault (2006), “Estimation of Objective and Risk-Neutral Distributions Based on Moments of Integrated Volatility,” Working Paper, CRDE, Université de Montréal.
- Gordon, Stephen and Pascal St-Amour (2004), “Asset Returns and State-Dependent Risk Preferences,” *Journal of Business and Economic Statistics*, vol. 22, 241–252.

- Goyal, Amit and Ivo Welch (2003), “Predicting the Equity Premium with Dividend Ratios,” *Management Science*, vol. 49, 639–654.
- Goyal, Amit and Ivo Welch (2005), “A Comprehensive Look at the Empirical Performance of Equity Premium Prediction,” NBER Working Paper No.10483.
- Guo, Hui and Robert F. Whitelaw (2006), “Uncovering the Risk-Return Relation in the Stock Market,” *Journal of Finance*, vol. 61, 1433–1463.
- Heston, Steven (1993), “A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options,” *Review of Financial Studies*, vol. 6, 327–343.
- Hodrick, Robert J. (1992), “Dividend Yields and Expected Stock Returns: Alternative Procedures for Inference and Measurement,” *Review of Financial Studies*, vol. 5, 357–386.
- Hsieh, David A. (1991), “Chaos and Nonlinear Dynamics: Application to Financial Markets,” *Journal of Finance*, vol. 46, 1839–1877.
- Jackwerth, Jens Carsten (2000), “Recovering Risk Aversion from Option Prices and Realized Returns,” *Review of Financial Studies*, vol. 13, 433–451.
- Jiang, George and Yisong Tian (2005), “Model-Free Implied Volatility and Its Information Content,” *Review of Financial Studies*, vol. 18, 1305–1342.
- Keim, Donald B. and Robert F. Stambaugh (1986), “Predicting Returns in the Stock and Bond Markets,” *Journal of Financial Economics*, vol. 17, 357–390.
- Lamont, Owen (1998), “Earnings and Expected Returns,” *Journal of Finance*, vol. 53, 1563–1587.
- Lettau, Martin and Sydney Ludvigson (2001), “Consumption, Aggregate Wealth, and Expected Stock Returns,” *Journal of Finance*, vol. 56, 815–849.
- Lewellen, Jonathan (2004), “Predicting Returns with Financial Ratios,” *Journal of Financial Economics*, vol. 74.
- Lewellen, Jonathan and Stefan Nagel (2006), “The Conditional CAPM Does Not Explain Asset-Pricing Anomalies,” *Journal of Financial Economics*, Forthcoming.
- Lewellen, Jonathan, Stefan Nagel, and Jay Shanken (2006), “A Skeptical Appraisal of Asset-Pricing Tests,” *Working Paper*, Dartmouth College.

- Lo, Andrew W. and A. Craig MacKinlay (1988), “Stock Market Prices do Not Follow Random Walks: Evidence from a Simple Specification Test,” *Review of Financial Studies*, vol. 1, 41–66.
- Meddahi, Nour (2002), “Theoretical Comparison Between Integrated and Realized Volatility,” *Journal of Applied Econometrics*, vol. 17, 479–508.
- Merton, Robert C. (1980), “On Estimating the Expected Return on the Market,” *Journal of Financial Economics*, vol. 8, 323–361.
- Mixon, Scott (2007), “The Implied Volatility Term Structure of Stock Index Options,” *Journal of Empirical Finance*, vol. 14, 333–354.
- Newey, Whitney K. and Kenneth D. West (1987), “A Simple Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix,” *Econometrica*, vol. 55, 703–708.
- Rosenberg, Joshua V. and Robert F. Engle (2002), “Empirical Pricing Kernels,” *Journal of Financial Economics*, vol. 64, 341–372.
- Schwert, G. William (1990), “Stock Volatility and the Crash of ’87,” *Review of Financial Studies*, vol. 3, 77–102.
- Shanken, Jay (1992), “On the Estimation of Beta-Pricing Models,” *Review of Financial Studies*, vol. 5, 1–33.
- Shiller, Robert J. (1984), “Stock Prices and Social Dynamics,” *Brookings Papers on Economic Activity*, vol. 2, 457–498.
- Stambaugh, Robert F. (1999), “Predictive Regressions,” *Journal of Financial Economics*, vol. 54, 375–421.
- Tauchen, George (2005), “Stochastic Volatility in General Equilibrium,” *Working Paper*.
- Todorov, Viktor (2006), “Pricing Diffusive and Jump Risk: What Can We Learn from the Variance Risk Premium,” Working Paper, Department of Economics, Duke University.
- Valkanov, Rossen (2003), “Long-Horizon Regressions: Theoretical Results and Applications,” *Journal of Financial Economics*, vol. 68, 201–232.
- Wu, Liuren (2005), “Variance Dynamics: Joint Evidence from Options and High-Frequency Returns,” Working Paper, Baruch College at CUNY.

Table 1 Summary Statistics for Key Variables

The sample period extends from 1990Q1 to 2005Q1. All variables are reported in annualized percentage form whenever appropriate. $R_{mt} - R_{ft}$ denotes the logarithmic return on the S&P500 in excess of the 3-month T-bill rate. IV_t denotes the “model-free” implied variance, or VIX index. RV_t refers to the “model-free” realized variance constructed from high-frequency 5-minute returns. The standard predictor variables include the price-earning ratio, $\log(P_t/E_t)$, the price-dividend ratio, $\log(P_t/D_t)$, the default spread (defined as the difference between Moody’s BAA and AAA bond yield indices), $DFSP_t$, the term spread (defined as the difference between the 10-year and 3-month Treasury yields), $TMSP_t$, and the stochastically de-trended risk free rate (defined as the 1-month T-bill rate minus its trailing twelve month moving averages), $RREL_t$. The consumption-wealth ratio, or CAY_t , is defined in Lettau and Ludvigson (2001), and the data is downloaded from their website.

	$R_{mt} - R_{ft}$	$IV_t - RV_t$	IV_t	RV_t	$\log(P_t/E_t)$	$\log(P_t/D_t)$	$DFSP_t$	$TMSP_t$	$RREL_t$	CAY_t
Summary Statistics										
Mean	5.97	21.40	36.30	14.90	3.16	3.91	0.83	1.90	-0.19	0.30
Std Dev	31.44	14.86	25.78	14.92	0.26	0.36	0.22	1.11	0.74	1.92
Skewness	-0.63	1.75	2.01	1.99	0.20	-0.10	0.95	-0.15	-0.40	-0.37
Kurtosis	3.49	6.38	8.14	6.52	2.45	1.70	2.94	1.90	3.15	2.08
AR(1)	-0.07	0.31	0.41	0.44	0.91	0.96	0.83	0.88	0.81	0.90
Correlation Matrix										
$R_{mt} - R_{ft}$	1.00	-0.28	-0.49	-0.57	-0.07	-0.03	-0.12	0.03	0.04	-0.06
$IV_t - RV_t$		1.00	0.86	0.50	0.17	0.09	0.35	-0.17	-0.21	-0.01
IV_t			1.00	0.87	0.37	0.35	0.35	-0.28	-0.24	-0.22
RV_t				1.00	0.47	0.52	0.25	-0.32	-0.20	-0.37
$\log(P_t/E_t)$					1.00	0.72	0.37	0.14	-0.24	-0.65
$\log(P_t/D_t)$						1.00	-0.00	-0.37	0.08	-0.89
$DFSP_t$							1.00	0.33	-0.56	0.04
$TMSP_t$								1.00	-0.31	0.39
$RREL_t$									1.00	-0.12
CAY_t										1.00

Table 2 Forecast Regressions for Quarterly Stock Market Return

The sample period extends from 1990Q1 to 2005Q1. Newey-West robust t -statistics with four lags are reported in parentheses. All variable definitions are identical to Table 1.

Regressors	Simple									Multiple				
Intercept (t -stat)	-12.19 (-1.55)	-6.13 (-0.83)	3.89 (0.75)	113.01 (2.27)	77.45 (1.84)	23.00 (1.26)	7.73 (0.69)	7.33 (1.70)	5.06 (1.11)	120.85 (1.90)	-13.89 (-2.18)	81.50 (1.12)	107.55 (1.27)	108.28 (1.72)
$IV_t - RV_t$ (t -stat)	0.86 (3.94)									0.98 (4.87)	0.87 (3.85)		0.97 (4.58)	1.11 (5.27)
IV_t (t -stat)		0.34 (2.38)												
RV_t (t -stat)			0.17 (0.56)											
$\log(P_t/E_t)$ (t -stat)				-2.81 (-2.08)						-3.57 (-2.09)		-1.99 (-1.04)	-3.22 (-1.42)	-3.51 (-2.20)
$\log(P_t/D_t)$ (t -stat)					-1.51 (-1.62)									
$DFSP_t$ (t -stat)						-1.67 (-0.87)								
$TMSP_t$ (t -stat)							-0.70 (-0.14)							4.75 (1.34)
$RREL_t$ (t -stat)								4.40 (0.91)						7.59 (2.29)
CAY_t (t -stat)									4.16 (2.29)		4.32 (2.58)	1.98 (0.82)	0.83 (0.37)	
Adj. R^2 (%)	15.14	6.32	-1.05	6.22	2.76	0.27	-1.63	-0.66	4.83	26.37	20.86	5.54	25.23	27.67

Table 3 Robustness Check with Alternative Variance Measures

The sample period extends from 1990Q1 to 2005Q1. Newey-West robust t -statistics with four lags are reported in parentheses. All variable definitions are identical to Table 1, except for IV_t^* which denotes the Black-Scholes implied variance, and RV_t^* which denotes the realized variance constructed from daily returns.

Regressors	Simple					Multiple $IV_t^* - RV_t$			Multiple $IV_t - RV_t^*$			Multiple $IV_t^* - RV_t^*$		
Constant (t -stat)	-7.11 (-0.97)	-5.35 (-0.92)	-3.35 (-0.56)	-2.75 (-0.42)	1.57 (0.27)	141.62 (2.27)	124.46 (1.42)	134.91 (2.16)	107.74 (1.79)	95.91 (1.18)	99.10 (1.67)	129.42 (2.15)	115.02 (1.35)	122.71 (2.03)
$IV_t^* - RV_t$ (t -stat)	0.50 (2.43)					0.71 (3.79)								
$IV_t - RV_t^*$ (t -stat)						0.70 (3.53)			0.78 (4.31)			0.76 (3.52)		
$IV_t^* - RV_t^*$ (t -stat)									0.82 (4.30)			0.64 (3.49)		
IV_t^* (t -stat)												0.62 (3.05)		
RV_t^* (t -stat)												0.71 (3.69)		
IV_t^* (t -stat)						0.22 (1.72)								
RV_t^* (t -stat)												0.23 (1.13)		
$\log(P_t/E_t)$ (t -stat)						-4.06 (-2.39)			-2.99 (-1.89)			-3.59 (-2.23)		
$TMSp_t$ (t -stat)						-3.61 (-1.53)			-2.68 (-1.24)			-3.22 (-1.42)		
$RREL_t$ (t -stat)						-4.22 (-2.62)			-2.85 (-1.90)			-3.59 (-2.32)		
CAY_t (t -stat)						5.61 (1.22)			1.85 (0.48)			3.25 (0.71)		
						6.78 (2.01)			4.66 (1.45)			4.76 (1.48)		
						1.06 (0.41)			0.75 (0.31)			0.88 (0.34)		
Adj. R^2 (%)	7.30	8.08	3.62	3.05	0.68	21.04	19.89	22.39	15.62	14.26	13.77	14.35	13.02	13.03

Table 4 Robustness Check with Monthly Stock Market Return

The sample period extends from January 1990 to January 2005. Newey-West robust t -statistics with twelve lags are reported in parentheses. All variable definitions are identical to Table 1, except being monthly observations.

Regressors	Simple								Multiple	
Intercept (t -stat)	-1.62 (-0.32)	-0.79 (-0.14)	5.07 (1.22)	103.66 (2.32)	74.91 (2.00)	8.95 (0.56)	8.07 (0.90)	6.81 (1.91)	105.79 (2.16)	102.22 (2.11)
$IV_t - RV_t$ (t -stat)	0.42 (2.87)								0.49 (2.68)	0.58 (3.03)
IV_t (t -stat)	0.21 (1.66)									
RV_t (t -stat)	0.10 (0.41)									
$\log(P_t/E_t)$ (t -stat)	-2.56 (-2.12)								-2.87 (-2.12)	-2.97 (-2.29)
$\log(P_t/D_t)$ (t -stat)	-1.45 (-1.75)									
$DFSP_t$ (t -stat)	0.03 (0.15)									
$TMSP_t$ (t -stat)	-0.76 (-0.19)									3.12 (0.85)
$RREL_t$ (t -stat)	1.09 (0.27)									2.69 (0.75)
Adj. R^2 (%)	1.24	0.58	-0.42	2.10	1.07	-0.52	-0.50	-0.50	3.98	3.37

Table 5 Robustness Check with Volatility Risk Premium

The sample period extends from 1990Q1 to 2005Q1. Newey-West robust t -statistics with four lags are reported in parentheses. All variable definitions are identical to Table 1, except for $\sqrt{IV_t}$ and $\sqrt{RV_t}$ which refer to the implied and realized volatilities, respectively.

Regressors	Simple			Multiple			
Constant (t -stat)	-27.81 (-2.84)	-15.82 (-1.17)	5.20 (0.61)	61.41 (1.03)	-25.97 (-2.87)	55.91 (0.73)	47.63 (0.83)
$\sqrt{IV_t} - \sqrt{RV_t}$ (t -stat)	4.42 (4.90)			4.15 (4.83)	4.07 (4.26)	4.12 (4.44)	4.39 (5.36)
$\sqrt{IV_t}$ (t -stat)		1.11 (1.75)					
$\sqrt{RV_t}$ (t -stat)			0.10 (0.13)				
$\log(P_t/E_t)$ (t -stat)				-2.29 (-1.45)		-2.14 (-1.04)	-2.02 (-1.39)
$TMSP_t$ (t -stat)							1.48 (0.42)
$RREL_t$ (t -stat)							6.01 (1.66)
CAY_t (t -stat)					2.72 (1.57)	0.37 (0.17)	
Adj. R^2 (%)	18.50	3.76	-1.66	22.38	19.81	21.05	21.32

Table 6 Robustness Check with Alternative Predictor Variables

The sample period extends from 1990Q1 to 2005Q1. Newey-West robust t -statistics with four lags are reported in parentheses. All variable definitions are identical to Table 1, except for $\log(D_t/E_t)$ which denotes the log payout ratio, λ_t which refers to the extracted volatility risk premium from Bollerslev, Gibson, and Zhou (2006), and RQ_t which denotes the realized quarticity as defined in the main text.

Regressors	Simple				Multiple			
Constant (t -stat)	6.85 (1.41)	7.66 (0.54)	-11.32 (-0.90)	3.74 (0.76)	129.03 (2.05)	132.27 (2.14)	119.75 (1.96)	118.26 (1.96)
$IV_t - RV_t$ (t -stat)					0.88 (5.02)	1.01 (4.92)	1.14 (5.71)	0.98 (5.78)
$\log(P_t/E_t)$ (t -stat)					-3.77 (-2.22)	-3.81 (-2.36)	-3.76 (-2.53)	-3.84 (-2.46)
$TMSP_t$ (t -stat)							4.91 (1.39)	6.41 (1.55)
$RREL_t$ (t -stat)							7.29 (2.36)	9.91 (2.70)
$R_{mt} - R_{ft}$ (t -stat)	-0.07 (-0.68)							
$\log(D_t/E_t)$ (t -stat)		1.68 (0.08)						
λ_t (t -stat)			0.84 (1.66)			-0.14 (-0.29)	-0.14 (-0.30)	
RQ_t (t -stat)				0.80 (2.92)	0.47 (2.22)			0.79 (2.33)
Adj. R^2 (%)	-1.17	-1.68	4.47	3.28	26.52	25.20	26.46	30.18

Table 7 Risk-Factor Test with Fama-French Portfolios

The regressions are based on monthly observations. The sample period extends from January 1990 to January 2005. The top panel reports the time-series regression coefficients for the twenty-five Fama and French (1993) portfolios on the market return and the variance difference as risk factors, along with the corresponding Newey-West robust t -statistics. The lower panel reports the cross-sectional regressions for the risk-premium estimates, with the t -statistics adjusted for heteroscedasticity, serial correlation, and the generated regressor problem following Shanken (1992).

Time-Series	Market					$IV_t - RV_t$				
Loading β	Small	2	3	4	Large	Small	2	3	4	Large
Low	1.47	1.46	1.37	1.30	1.03	-1.00	-1.10	-1.05	-1.01	-0.80
2	1.15	1.04	1.02	0.94	0.89	-0.99	-0.94	-1.11	-1.15	-1.08
3	0.90	0.85	0.84	0.89	0.79	-0.82	-0.92	-1.01	-1.14	-0.99
4	0.79	0.82	0.79	0.81	0.68	-0.81	-1.00	-0.93	-0.86	-0.86
High	0.84	0.90	0.85	0.81	0.79	-0.94	-0.89	-0.96	-0.84	-1.05
t -statistic	Small	2	3	4	Large	Small	2	3	4	Large
Low	14.43	19.50	20.47	25.54	33.98	-1.38	-1.70	-1.76	-2.05	-2.21
2	12.21	15.30	20.29	17.56	22.52	-1.67	-1.76	-2.47	-2.71	-2.93
3	11.97	13.14	14.20	14.08	15.78	-1.62	-2.08	-2.33	-2.74	-2.64
4	10.66	11.57	11.85	12.84	11.22	-1.72	-2.23	-2.31	-2.29	-2.78
High	10.98	11.16	10.75	10.71	10.48	-2.06	-1.94	-2.28	-2.13	-3.09
Adj. R^2 (%)	Small	2	3	4	Large	Small	2	3	4	Large
Low	49.38	64.28	67.07	77.00	88.02	1.60	2.95	3.17	3.81	4.46
2	44.53	63.28	73.64	76.86	78.57	2.57	4.30	7.65	10.30	10.30
3	48.31	59.96	67.00	66.54	67.36	3.27	6.14	8.52	9.82	9.35
4	45.44	55.22	55.93	64.52	50.05	3.95	7.18	6.84	6.32	6.92
High	46.25	52.60	55.12	54.06	45.92	4.94	4.25	6.16	5.02	7.16
Cross-Section	Market					$IV_t - RV_t$				
Risk Premium λ	10.12					-10.73				
t -statistic	2.45					-1.54				




Figure 1 Implied Variance, Realized Variance, and Variance Risk Premium

This figure plots the implied variance (top panel, blue line), the realized variance (middle panel, red line), and the difference (bottom panel, black line) for the S&P 500 market index from 1990 Q1 to 2005 Q1. The shaded areas represent NBER recessions.




Figure 2 Candidate Return Predictors and Excess Stock Returns

The figure plots the quarterly S&P 500 excess returns (thin red line) together with (lagged) predictor variables (blue lines) from 1990 Q1 to 2005 Q1. The top panel illustrates the variance difference, or variance risk premium, the middle panel graphs the P/E ratio, and the bottom panel depicts the consumption-wealth ratio CAY.




Figure 3 Variance Risk Premium and GDP Growth

The figure plots the GDP growth rates (thin red line) together with the lagged variance differences (blue line), or variance risk premia, from 1990 Q1 to 2005 Q1. Both of the series are standardized to have mean zero and variance one.