UvA-DARE (Digital Academic Repository)

The pricing of long and short run variance and correlation risk in stock returns

Cosemans, M.

Publication date

2011

Document Version

Submitted manuscript

[Link to publication](#)

Citation for published version (APA):

Cosemans, M. (2011). *The pricing of long and short run variance and correlation risk in stock returns*. University of Amsterdam Business School.
http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1825934

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

The Pricing of Long and Short Run Variance and Correlation Risk in Stock Returns

Mathijs Cosemans

University of Amsterdam

April 29, 2011

Abstract

This paper studies the pricing of long and short run variance and correlation risk. The predictive power of the market variance risk premium for returns is driven by the correlation risk premium and the systematic part of individual variance premia. Furthermore, I find that aggregate volatility risk is priced in the cross-section because shocks to average stock volatility and correlation are priced. Both long and short run volatility and correlation factors have explanatory power for returns. Finally, I resolve the idiosyncratic volatility puzzle by showing that short-term idiosyncratic risk is positively priced whereas long-term idiosyncratic volatility carries a negative price.

Keywords: variance risk, correlation risk, idiosyncratic risk, return predictability

JEL classification: G12, G14

Contact details: Mathijs Cosemans, University of Amsterdam Business School, Roetersstraat 11, 1018 WB Amsterdam, Netherlands, E-mail: M.Cosemans@uva.nl. I thank Andrew Ang, Dion Bongaerts, Riccardo Colacito, Mathijs van Dijk, Joost Driessen, Robert Engle, Rik Frehen, Andrew Patton, Jose Gonzalo Rangel, Peter Schotman, Kevin Sheppard and seminar participants at RSM Erasmus University, Maastricht University, Ortec Finance, the 2010 Annual Meeting of the Society for Financial Econometrics (SoFiE), and the 2011 Annual Winter Meetings of the Econometric Society for helpful comments and suggestions. Part of this paper was written while I was a visiting research fellow at Columbia University.

1 Introduction

Traditionally, risk in financial markets is measured as the risk that returns vary over time. However, a large body of empirical evidence shows that risk itself also changes over time.¹ During the credit crisis of 2008, for example, implied market volatility increased from 20% in August to 80% two months later. Time-varying market risk can be a priced risk factor in intertemporal models because it affects the future risk-return tradeoff. Since changes in market volatility reflect variation in market-wide correlations and fluctuations in average stock volatility, individual volatility risk and correlation risk should be priced if aggregate volatility risk is priced. Furthermore, macroeconomic and firm-specific factors induce high- and low-frequency movements in volatilities and correlations.² Because shocks to short-term variances and correlations are transitory, they can be priced differently from innovations in long-term variances and correlations that are more permanent.

In this paper I analyze the pricing of long and short run variance and correlation risk in aggregate returns and in the cross-section of individual stock returns. My first contribution is to study whether the correlation risk premium and individual variance risk premium predict stock market returns. Bollerslev, Tauchen, and Zhou (2009) show that the market variance risk premium, measured as the difference between implied and realized market variance, predicts the equity premium. Recent work by Driessen, Maenhout, and Vilkov (2009) shows that market variance risk is priced in option markets only because of priced correlation risk, not individual variance risk.

Using data for S&P 100 index options and for options on the S&P 100 components, I show that the predictive power of the market variance premium is driven by the corre-

¹Surveys of the vast literature on time-varying volatility are given by Bollerslev, Engle, and Nelson (1994) and Ghysels, Harvey, and Renault (1996). Empirical evidence that correlations vary over time and tend to increase when stock prices decrease is presented in Bollerslev, Engle, and Wooldridge (1988) and Moskowitz (2003).

²Engle and Lee (1999) and Chernov, Gallant, Ghysels, and Tauchen (2003) find that long run and short run component models better explain equity volatility and Engle and Rangel (2008) link variation in long run market volatility to macroeconomic variables. Campbell, Lettau, Malkiel, and Xu (2001) document an upward trend in idiosyncratic volatility and a downward trend in correlations until 1997. However, Bekaert, Hodrick, and Zhang (2010) find no evidence of such a trend when extending the sample till 2008.

lation risk premium and the systematic component of the average variance risk premium in individual stock options. In contrast, idiosyncratic variance risk premia do not predict market returns. These results hold for both monthly and quarterly forecasting horizons and are robust to the inclusion of traditional return predictors and the impact of the recent financial crisis.

My second goal is to examine whether long and short run correlation risk and average volatility risk are priced in the cross-section. Ang, Hodrick, Xing, and Zhang (2006) document a negative relation between returns and innovations in market volatility. Adrian and Rosenberg (2008) decompose market volatility into long- and short-term components and find that shocks to both parts carry a significant price of risk. Intertemporal models predict that the price of risk on these factors should be negative since risk-averse investors want to hedge against a sudden increase in aggregate uncertainty and a loss of diversification benefits. Consequently, stocks that pay off when investment opportunities deteriorate should have lower expected returns.

I study these cross-sectional predictions by decomposing volatilities and correlations into high- and low-frequency components using the Spline-GARCH model proposed by Engle and Rangel (2008) for volatilities and extended by Rangel and Engle (2009) for correlations. This approach models the short-term component of volatility as a GARCH process and the long-term component using a quadratic spline. High- and low-frequency volatilities are combined with a factor model to capture long and short run patterns in correlations. The factor model allows for a parsimonious representation of the correlation structure and makes it possible to distinguish between cycles in aggregate volatility and patterns in idiosyncratic volatility.

Using a sample of individual stocks, I document large movements in both long and short run components of market volatility, idiosyncratic volatility, and correlations. Furthermore, I identify important differences in the cyclical patterns in idiosyncratic volatility across stocks. These risk dynamics generate significant premia in the cross-section of returns. In particular, I find that market volatility risk is priced because innovations

in both average stock volatility and market-wide correlations are priced, in line with the predictions of the ICAPM. Consistent with the results from the time series analysis, I find that the systematic component of average stock volatility is a priced risk factor. In contrast, shocks to average idiosyncratic volatility are not priced. In addition, I document that both long-term and short-term volatility and correlation factors carry a negative price of risk. The long and short run factors remain significant when jointly included in the ICAPM, which suggests that they capture different sources of risk. The explanatory power of the variance and correlation risk factors is also not subsumed by traditional size, value, momentum, and liquidity factors.

My third objective is to study the pricing of long-term and short-term components of idiosyncratic risk.³ Portfolio theory and the CAPM predict that idiosyncratic risk should not be priced because it can be diversified away. However, Ang, Hodrick, Xing, and Zhang (2006) find a negative cross-sectional relation between idiosyncratic volatility and average stock returns. In contrast, Fu (2009) finds a positive relation between idiosyncratic volatility and returns when using a conditional measure of idiosyncratic risk.

I find that the short-term component of idiosyncratic volatility is positively priced while the long-term component carries a negative price of risk. Most importantly, both components continue to be priced when simultaneously added to the model. Thus, short run idiosyncratic risk is priced differently from long run idiosyncratic volatility, which reconciles the conflicting evidence in the literature. The pricing of the short-term, conditional measure of idiosyncratic risk is explained by exposure to volatility risk. In contrast, the pricing of long-term, unconditional idiosyncratic risk is robust to the inclusion of volatility and correlation risk factors.

The paper proceeds as follows. Section 2 reviews literature on the pricing of variance and correlation risk in aggregate returns and in the cross-section. Section 3 studies the predictive power of variance and correlation premia for market returns. Section 4 describes the methodology used to measure long and short run volatilities and correlations.

³This analysis focuses on the pricing of the *level* of idiosyncratic volatility while the discussion in the previous paragraph is about the pricing of sudden *changes* in aggregate and average stock volatility.

In Section 5 I examine the cross-sectional pricing of shocks to these volatility and correlation components and study the relation between long and short run idiosyncratic risk and returns. Section 6 concludes.

2 Literature on Variance and Correlation Risk

2.1 Variance and Correlation Premia in Aggregate Returns

Bollerslev, Tauchen, and Zhou (2009) show that the market variance risk premium has strong predictive power for aggregate returns, particularly at intermediate horizons. This result continues to hold after controlling for traditional return predictors like the dividend yield and the consumption-wealth ratio. They set up a general equilibrium model to explain the predictive power of the variance premium that extends the long run risk model of Bansal and Yaron (2004) by allowing for time-varying volatility of consumption growth volatility.⁴ The equity premium in their model consists of the standard consumption risk term and a volatility risk term. The authors argue that the variance risk premium isolates this volatility of volatility factor, which means that movements in the variance risk premium reflect changes in variance risk.

An alternative source of variation in the variance risk premium is time-varying risk aversion. The long run risk model assumes that risk aversion is constant but Campbell and Cochrane (1999) argue that risk aversion varies over time due to habit formation. Bollerslev, Gibson, and Zhou (2009) derive an approximate analytic relation between the variance risk premium and relative risk aversion of the representative investor and link time variation in the premium to macroeconomic variables. Todorov (2009) argues that the variance risk premium is a compensation for stochastic volatility and price jumps. He documents that even though the effect of jumps on volatility disappears fast, their effect

⁴Drechsler and Yaron (2011) extend the long run risk model by allowing for infrequent jumps in long run consumption growth rates and in the volatility of consumption. They argue that changes in the variance premium reflect time variation in agents' perceptions of the risk of these big shocks to the state of the economy.

on the variance risk premium is long-lasting due to a time-varying attitude of investors towards jumps. Bekaert and Engstrom (2010) set up a consumption based model with habit-based preferences as in Campbell and Cochrane (1999) but with non-linearities in consumption growth. In their framework the variance premium increases with shocks that generate negative skewness of consumption.

Driessen, Maenhout, and Vilkov (2009) decompose the market variance risk premium into a correlation risk premium and individual variance risk premia using index options and individual stock options. They document a large variance premium in index options but find no evidence of a variance premium in individual options. Given the decomposition of stock market volatility, they interpret this as indirect evidence of priced correlation risk. According to this explanation, index options are more expensive than individual options because they can be used to hedge against an increase in market-wide correlations that reduces diversification benefits.

If market variance risk is priced because of correlation risk, then a testable implication of the theoretical framework of Bollerslev, Tauchen, and Zhou (2009) is that the predictive power of the aggregate variance premium is driven by the correlation premium. Pollet and Wilson (2010) emphasize another channel through which correlations can predict stock market returns that is motivated by the Roll critique. They argue that changes in market-wide correlations reveal changes in true aggregate risk because correlations are driven by the exposure of stocks to the return on the latent aggregate wealth portfolio. In contrast to Goyal and Santa-Clara (2003), they find that average stock variance has no predictive power for aggregate returns. Guo and Savickas (2008) show that value-weighted average idiosyncratic volatility does have significant predictive power for aggregate returns when combined with stock market volatility.

2.2 Cross-Sectional Evidence on Variance and Correlation Risk

Ang, Hodrick, Xing, and Zhang (2006) show that innovations in aggregate volatility carry a significantly negative price in the cross-section of stock returns. They argue that

investors want to hedge against increases in market volatility because periods of high volatility tend to coincide with a market decline. The resulting hedging demand for stocks with a positive exposure to changes in market volatility increases their price and lowers their expected returns. Carr and Wu (2009) find large cross-sectional variation in the variance risk premium on individual stock options. They show that the premium is not explained by exposure to traditional risk factors and conclude that it is driven by an independent variance risk factor. Adrian and Rosenberg (2008) specify a two-component volatility model and find that both long run and short run aggregate volatility risk is negatively priced. They argue that the short-term volatility component captures market skewness risk, which they interpret as a measure of the tightness of financial constraints. The long-term component is associated with business cycle risks.

The decomposition of market variance into individual variances and correlations implies that if market variance risk is a priced factor in the cross-section of returns, then correlation risk, individual volatility risk or both should be priced too. Krishnan, Petkova, and Ritchken (2009) present evidence of a negative correlation risk premium in stock markets, after controlling for asset volatility and other risk factors. In contrast, Chen and Petkova (2010) find that only innovations to average stock variance are priced in a cross-section of portfolios sorted on idiosyncratic volatility. By decomposing stock variances into systematic and idiosyncratic parts, they show that only the systematic component of average variances is priced. In a recent paper, Schurhoff and Ziegler (2010) also decompose variance risk premia in options into systematic and idiosyncratic components. They find that systematic variance risk is negatively priced in the cross-section of options whereas idiosyncratic variance risk carries a positive price.

The above papers study the relation between returns and sudden *changes* in variances and correlations, motivated by an intertemporal asset pricing model. Another strand of literature studies the relation between the *level* of idiosyncratic volatility and stock returns. The CAPM predicts that idiosyncratic risk should not be priced because it can be diversified away. However, Ang, Hodrick, Xing, and Zhang (2006) find a negative

cross-sectional relation between idiosyncratic risk and expected returns. In contrast, Fu (2009) finds a positive relation between idiosyncratic volatility and returns when using a conditional measure of idiosyncratic risk. This result is consistent with the theoretical model of Merton (1987) in which idiosyncratic risk is positively priced in information-segmented markets if investors do not hold well-diversified portfolios. Chen and Petkova (2010) argue that the idiosyncratic volatility puzzle of Ang, Hodrick, Xing, and Zhang (2006) can be explained by the different exposures of high and low idiosyncratic volatility portfolios to innovations in average stock variance.

I contribute to the time series literature by first decomposing the aggregate variance risk premium into the correlation risk premium and the average variance risk premium in individual options. Subsequently, I study whether the predictive power of the market variance premium is driven by the correlation premium, the average variance premium, or both. I extend the cross-sectional work by decomposing both volatilities and correlations into long- and short-term components. Subsequently, I study the cross-sectional pricing of each component to determine whether aggregate volatility risk is priced because of correlation risk or individual volatility risk. Finally, I provide evidence on the pricing of long and short run components of idiosyncratic risk.

3 Forecasting Stock Market Returns with Variance and Correlation Premia

3.1 Measurement of Variance and Correlation Risk Premia

The variance risk premium is formally defined as the difference between the risk-neutral and physical expectation of future variance. Britten-Jones and Neuberger (2000) demonstrate that the risk-neutral expectation of integrated variance is equal to the model-free implied variance under the assumption that the underlying asset price is continuous but volatility is stochastic. Specifically, they define the model-free implied variance of asset i

over the interval Δs as

$$IV_{is} = 2 \int_0^\infty \frac{C_{is}(s + \Delta s, K) - C_{is}(s, K)}{K^2} dK, \quad (1)$$

where $C_{is}(T, K)$ is the price on day s of a European call option on asset i with strike price K and maturity at time T . I measure the implied variance for a fixed 30-day maturity, i.e., Δs equals 30, because the options needed to compute the implied variance are most liquid for short maturities. Moreover, the VIX index that is often used by investors as a measure of implied volatility is also defined for a 30-day maturity. I construct the implied variance for each month t , IV_{it} , by taking the average of all daily implied variances within the month, which reduces the noise in the daily measures. Because implied variance is computed directly from option prices, no option pricing model is needed. In contrast, implied variance backed out from the Black and Scholes (1973) model is a flawed estimate of risk-neutral expected variance because the model assumes that volatility is constant. Although the model-free implied variance in equation (1) is defined as the integral over a continuum of strike prices, in practice only a finite number of different strikes is available. However, Jiang and Tian (2005) show that discretization errors are small when the integral is calculated from a limited number of options.

Following Bollerslev, Tauchen, and Zhou (2009), I also use a model-free measure of realized variance. In particular, realized variance for asset i over the interval Δt is calculated by summing squared intraday returns

$$RV_{it} = \sum_{j=1}^n \left[p_{it-1+\frac{j}{n}(\Delta)} - p_{it-1+\frac{j-1}{n}(\Delta)} \right]^2, \quad (2)$$

where p_{it} denotes the logarithmic price of asset i and n is the number of price observations within the interval Δt .⁵ Andersen, Bollerslev, Diebold, and Ebens (2001) show that this

⁵Since realized variance in (2) is measured ex-post, it is strictly speaking not equal to the physical expectation of future variance. However, it has the advantage that it is directly observable, which is important for forecasting purposes. An alternative approach is followed by Bollerslev, Gibson, and Zhou (2009), who rely on a stochastic volatility model to obtain a forecast of future return variation. Hence, their realized variance measure is no longer model-free and can be affected by model misspecification.

approach provides a more precise measure of realized variance than traditional measures of return variation based on returns sampled at a lower frequency. In practice, market microstructure frictions, such as the bid-ask bounce and price discreteness, put an upper limit on the data frequency that can be used to estimate realized variance. While these issues are less important for the market index and for liquid stocks, they can have a serious impact on stocks that are less frequently traded. I therefore follow Driessen, Maenhout, and Vilkov (2009) and restrict the analysis to stocks that are included in the S&P 100 index because these are most actively traded. An additional benefit is that all stocks included in the S&P 100 have exchange-listed options, which are required to compute implied variances.

The aggregate and individual variance risk premium for a one-month horizon are defined as the difference between the measures of implied variance in (1) and realized variance in (2),

$$VRP_{Mt} = IV_{Mt} - RV_{Mt} \quad \text{and} \quad VRP_{it} = IV_{it} - RV_{it}. \quad (3)$$

I define the average individual variance premium, $\overline{VRP}_t$, as the value-weighted average of the variance premia on all stocks in the market index. Following Schurhoff and Ziegler (2010), I further decompose individual variance risk premia into systematic and idiosyncratic components by assuming a one-factor structure

$$VRP_{it} = \beta_i^2 VRP_{Mt} + VRP_{\epsilon it}, \quad (4)$$

where β_i is the stock's exposure to the excess market return.⁶

Estimates of implied and realized correlations are needed to construct the correlation risk premium. These measures can be derived by decomposing the market variance into

⁶Schurhoff and Ziegler (2010) find that assuming constant betas does not affect the results.

individual variances and correlations,

$$VAR_{Mt} = \sum_{i=1}^{N_t} w_{it}^2 VAR_{it} + \sum_{i=1}^{N_t} \sum_{j \neq i}^{N_t} w_{it} w_{jt} \sqrt{VAR_{it}} \sqrt{VAR_{jt}} COR_{ijt}, \quad (5)$$

where w_{it} is the index weight of stock i and N_t the number of stocks in the index at time t .⁷

Because I focus on market-wide correlation risk, I assume that all pairwise correlations are equal, i.e., $COR_{ijt} = COR_t \forall i, j, i \neq j$. Given this assumption and the decomposition of market variance in (5), average implied and realized correlation at time t can be calculated as⁸

$$IC_t = \frac{IV_{Mt} - \sum_{i=1}^{N_t} w_{it}^2 IV_{it}}{\sum_{i=1}^{N_t} \sum_{j \neq i}^{N_t} w_{it} w_{jt} \sqrt{IV_{it}} \sqrt{IV_{jt}}}, \quad (6)$$

$$RC_t = \frac{RV_{Mt} - \sum_{i=1}^{N_t} w_{it}^2 RV_{it}}{\sum_{i=1}^{N_t} \sum_{j \neq i}^{N_t} w_{it} w_{jt} \sqrt{RV_{it}} \sqrt{RV_{jt}}}. \quad (7)$$

The correlation risk premium is defined as the difference between the measures of implied correlation in (6) and realized correlation in (7),

$$CRP_t = IC_t - RC_t. \quad (8)$$

3.2 Data Description

Data for S&P 100 index options and for options on all stocks in the S&P 100 comes from OptionMetrics. The option data is daily and covers the period January 1996 to December 2008. The zero-coupon interest rate curve and underlying stock prices are retrieved from

⁷The S&P 100 is a value-weighted index that is rebalanced quarterly. Therefore I compute the weight of each firm based on its market capitalization.

⁸Another way to measure realized correlation is to compute the weighted average of rolling correlations between stocks. I measure realized correlation using (7) to be consistent with the calculation of implied correlation. The difference between this estimate of realized correlation and the average rolling correlation is negligible.

CRSP. Following Driessen, Maenhout, and Vilkov (2009), I exclude options with zero bid price, zero open interest, or missing delta. Furthermore, I remove calls with delta smaller than 0.15 and puts with delta larger than -0.05 because of their high implied volatility. Since a cross-section of liquid options is needed to calculate model-free implied variances, I focus on short-term options. As in Carr and Wu (2009), on each day and for each stock I choose the options with the two nearest maturities, except when the shortest maturity is within eight days. In that case I pick the next two maturities to avoid potential microstructure frictions of very short-term options. I only consider out-of-the-money and at-the-money options because these options tend to be more liquid than in-the-money options. Finally, I only calculate implied variance on a given day for stocks that have at least two calls and two puts available to minimize errors from discretization.

I translate option prices into implied volatilities using the Black-Scholes model. I then interpolate these implied volatilities using a cubic spline across moneyness levels, as proposed by Jiang and Tian (2005). To obtain implied volatilities for moneyness levels beyond the available range I extrapolate the implied volatilities at the highest and lowest strike price. This interpolation-extrapolation procedure generates a grid of 1000 implied volatilities for moneyness levels between 0.01 and 3.00. I convert the extracted implied volatilities back into call option prices using the Black-Scholes model and use these call prices to calculate the implied variance in equation (1). Finally, to obtain the implied variance for a fixed 30-day maturity, I linearly interpolate between the implied variances for the two maturities closest to the 30-day maturity.

To calculate the realized variance in equation (2), I follow Andersen, Bollerslev, Diebold, and Ebens (2001) and use transaction prices from the NYSE trades and quotes (TAQ) database for the components of the S&P 100 and prices from Tick Data for the S&P 100 index. I estimate realized variances using five-minute returns because this sampling frequency provides a reasonable balance between efficiency and robustness to

microstructure noise.⁹ For each day I construct equally spaced returns by computing the logarithmic difference in transaction prices at or immediately before each five-minute mark. To purge the returns from any negative autocorrelation due to the uneven spacing of the observed prices and the bid-ask bounce I de-mean and filter the raw returns using an MA(1) model. Each day I calculate 78 five-minute returns from 9:30am until 16:00pm, including the close-to-open overnight return. Hence, the realized variance estimate for a typical month with 22 trading days is based on $22 \times 78 = 1716$ returns.

Table 1 reports summary statistics for the monthly market return, market variance, average stock variance, and average correlation. The mean excess return on the S&P 100 index is 0.10% per month. The average implied market variance is 37.90 per month whereas the average realized market variance is only 26.50, which implies that the variance risk premium equals 11.40 (in percentages squared). In contrast, the value-weighted average implied variance of individual stocks is smaller than the average realized variance, which means that the average individual variance risk premium is negative. The average implied correlation between all stocks in the index is 0.36 and exceeds the average realized correlation of 0.21, which results in a correlation risk premium of 0.15. The positive premia for market variance risk and correlation risk and the negative premium for individual variance risk are consistent with results documented by Bollerslev, Tauchen, and Zhou (2009) and Driessen, Maenhout, and Vilkov (2009). Panel B shows that the contemporaneous correlation between the market variance premium and the premia for individual variance risk and correlation risk is positive. However, the correlation between the individual variance premium and correlation premium is close to zero, suggesting that these components capture orthogonal parts of the aggregate variance premium.

[Table 1 about here.]

Figure 1 plots the implied and realized market variance and the aggregate variance

⁹The choice of five-minute returns follows Andersen, Bollerslev, Diebold, and Ebens (2001). Driessen, Maenhout, and Vilkov (2009) sample returns at a daily frequency but Bollerslev, Gibson, and Zhou (2009) show that using realized volatilities based on daily returns results in inefficient estimates of the variance risk premium, which loses some of its predictive power for stock returns. I find that calculating realized variances using 15-minute returns instead of five-minute returns does not affect the forecasting results.

premium. Both variances and the risk premium increase sharply during periods of market turmoil, such as the Asian crisis in 1997, the LTCM and Russian crisis in 1998, the uncertainty surrounding the Iraq war in 2003, and the quant crisis in 2007. During the credit crisis in 2008 realized market variance rises above implied variance, which leads to a negative variance premium.

[Figure 1 about here.]

The average implied and realized variances of individual stocks depicted in Figure 2 also rise during these crisis periods. However, the increase in average realized variance is much larger than the increase in implied variance, which translates into a negative individual variance risk premium in most crisis periods. Further decomposing the variance risk premium on individual stocks into systematic and idiosyncratic parts using equation (4), I find that the systematic component is usually positive while the idiosyncratic variance premium is negative. This confirms the finding of Schurhoff and Ziegler (2010) that systematic and idiosyncratic variance risk premia have an opposite effect on the total variance premium on individual stocks.

[Figure 2 about here.]

Figure 3 shows that the correlation risk premium also exhibits strong time variation and increases during turbulent market conditions. Furthermore, unlike the variance premia the correlation risk premium remains positive during the 2008 financial crisis.

[Figure 3 about here.]

3.3 Predicting Returns with Variance and Correlation Premia

Table 2 shows results for monthly predictability regressions of the market return on lagged variance and correlation risk premia. All regressors are standardized to have mean zero and standard deviation equal to one for ease of interpretation. The first column in Panel

A of the table indicates that the finding of Bollerslev, Tauchen, and Zhou (2009) that the market variance premium predicts monthly stock market returns also holds when the sample includes the credit crisis in 2008. The relation between the market variance premium and the equity premium is positive and significant and the adjusted R^2 of the regression is 3.35%.

Decomposing the market variance premium into the average variance premium on individual stocks and the correlation premium reveals that both components have predictive power for market returns. The predictive power is not only statistically significant but also economically large. Specifically, a one standard deviation increase in either the average variance premium or correlation premium predicts an increase in the equity premium of about 0.90%. Consistent with the low correlation between the average variance premium and the correlation premium in Table 1, both predictors remain significant when jointly added to the regression. Also including the market variance premium does not increase forecasting power, which suggests that all information embedded in the aggregate variance premium is captured by the average variance premium and the correlation premium. Since Figure 2 shows that the idiosyncratic part of the average variance premium behaves very differently from the systematic part, I include these components in separate predictability regressions. Interestingly, the predictive power of the average variance premium is completely driven by the systematic part of individual variance premia. In contrast, the average idiosyncratic variance premium has no predictive power for market returns.

In Panel B I control for standard return predictors. Following Bollerslev, Tauchen, and Zhou (2009), I include the default spread, defined as the yield spread between Moody's Baa- and Aaa-rated corporate bonds, the log of the smoothed price-earnings ratio for the S&P 500, defined as the ratio of the price of the index divided by the twelve-month trailing moving average of aggregate earnings, the term spread, defined as the difference between the ten-year and three-month Treasury yield, and the consumption-wealth ratio

of Lettau and Ludvigson (2001).¹⁰ The coefficient on the market variance premium in the return predictability regressions remains economically and statistically significant after adding the control variables.¹¹ The adjusted R^2 of the one-month ahead predictability regression increases slightly after including the macroeconomic predictors. The average variance risk premium and the correlation risk premium also retain their predictive power in the presence of these traditional predictive variables.

[Table 2 about here.]

In Table 3 I study whether variance and correlation risk premia also have predictive power for stock market returns at a quarterly horizon. Because the quarterly returns are based on overlapping monthly observations, I report t -statistics based on Hodrick (1992) standard errors that are robust to heteroskedasticity and autocorrelation and account for the overlap in the regressions. Column 1 in Panel A shows that the market variance premium is also significant in these quarterly regressions. A one standard deviation increase in the aggregate variance premium corresponds to a 2.61% increase in the predicted quarterly market return. Both the average variance premium and the correlation premium also predict quarterly returns. The predictive power of the average variance premium derives from the information embedded in its systematic part, as the average idiosyncratic variance premium does not predict returns. The explanatory power of the variance and correlation premia in the quarterly return regressions is much higher than in the corresponding monthly regressions. However, the higher R^2 only reflects the persistence of the variance and correlation risk premia. In fact, scaling the quarterly returns by the horizon (i.e., dividing the regression coefficients by 3) reveals that the economic magnitude of the predictability is similar to that for the monthly returns in Table 2.

¹⁰Using other predictors like the dividend yield leads to very similar results. Moreover, in line with Bollerslev, Tauchen, and Zhou (2009) I find that controlling for the traditional risk-return tradeoff by including realized variances and correlations also does not explain the predictive power of variance and correlation risk premia.

¹¹Most traditional predictors are insignificant when jointly added to the regression due to collinearity. When included separately they do show up significantly and the signs of the coefficients are consistent with the literature.

Results reported in Panel B show that the market variance premium and the correlation risk premium remain significant predictors for quarterly returns after controlling for the standard predictive variables. However, the adjusted R^2 goes up considerably, reflecting the strong predictive power of the default spread and the consumption wealth ratio at the quarterly horizon. The average variance premium is no longer significant after including these traditional return predictors. Its predictive power is weaker than that of the correlation risk premium because the idiosyncratic part of the average variance premium does not contain information about future stock market returns, thereby offsetting the predictive power of the systematic component.

[Table 3 about here.]

Table 1 shows that the distribution of variance risk premia is far from normal while the correlation risk premium appears to be fairly normally distributed. The non-normality of variance premia is due to the huge increase in variances during the financial crisis at the end of 2008, as shown in Figures 1 and 2. As a robustness check, I therefore also report predictability results for the pre-crisis period from January 1996 to December 2007. Comparing these results reported in Table 4 to those in Tables 2 and 3 shows the influence of the outliers in 2008 on the predictive power of variance and correlation risk premia. Panel A shows that the market variance premium and the correlation premium are also significant predictors of stock market returns in the pre-crisis period. The predictive power of the aggregate variance premium as measured by the R^2 is smaller than for the full sample period but similar to the R^2 reported by Bollerslev, Tauchen, and Zhou (2009). The coefficient estimates are almost the same as those in Table 2, which indicates that the predictability is similar in economic terms. The coefficient on the average variance premium is insignificant, although its systematic component does have some predictive power. Results for a quarterly forecasting period in Panel B confirm that the predictive power of the market variance premium in the pre-crisis period is driven by the correlation risk premium and the systematic part of the average variance premium.

[Table 4 about here.]

Overall, the results from the return forecasting regressions show that the predictive power of the market variance risk premium is driven by the correlation premium and the systematic part of individual variance premia. Idiosyncratic variance premia do not forecast the equity premium. The predictive power of variance and correlation risk premia is even stronger when the sample includes the credit crisis in 2008. These results hold for both monthly and quarterly forecasting horizons and are robust to the inclusion of standard return predictors.

4 Long and Short Run Variances and Correlations

4.1 ICAPM with Time-Varying Variances and Correlations

The existence of variance and correlation risk premia in options indicates that investors are willing to pay a premium to insure against shocks to systematic volatility and market-wide correlations. In this section I study the pricing implications of investor aversion to variance and correlation risk for the cross-section of stocks. The theoretical motivation for a risk premium for variance and correlation risk is given by the intertemporal capital asset pricing model (ICAPM) of Merton (1973). The main premise of this model is that investors want to hedge against a deterioration of investment opportunities. Intuitively, assets that perform poorly when the market return is low, when expected future market returns fall, or when expected future market variances increase should have higher expected returns. Consequently, any variable that forecasts future market returns or variances is a relevant state variable of the pricing kernel.

Because market variance changes due to fluctuations in the variance of individual stocks and changes in correlations between these stocks, the ICAPM predicts that both shocks to average stock variance and innovations in market-wide correlations should be priced. In addition, several studies find that two component volatility specifications

better explain equity volatility than single-factor models because news can affect stock prices at different horizons (see e.g., Engle and Lee (1999), Chernov, Gallant, Ghysels, and Tauchen (2003), and Engle and Rangel (2008)). Colacito, Engle, and Ghysels (2009) and Rangel and Engle (2009) also identify long and short run movements in correlations. Shocks to the short-term component of variances and correlations are transitory while innovations in the long-term component are more permanent. Consequently, long and short run variance and correlation risk can carry a different price in the cross-section.

I test these predictions by estimating the following ICAPM specification,

$$r_{it+1} = \alpha_i + \gamma Cov_t(r_{it+1}, r_{Mt+1}) + \lambda' Cov_t(r_{it+1}, S_{t+1}) + \epsilon_{it+1}, \quad (9)$$

where γ is the relative risk aversion coefficient and where the elements in λ are the prices of hedge-related risks in S_{t+1} . The covariances in (9) measure the exposure of asset i to movements in the market return and to innovations in the state variables. Since my aim is to study the pricing of long and short run variance and correlation risk, I include in S_{t+1} shocks to the long- and short-term component of market volatility, average stock volatility, and average correlations.

4.2 Specification of Volatility and Correlation Components

I separate volatilities and correlations into long- and short-term components using the Spline-GARCH model first proposed by Engle and Rangel (2008) for volatilities and subsequently extended by Rangel and Engle (2009) for correlations.¹² This semi-parametric approach models the high-frequency component of aggregate and idiosyncratic volatility as a GARCH process and the low-frequency component using a quadratic spline.

¹²An alternative approach to split correlations into long- and short-term components is suggested by Colacito, Engle, and Ghysels (2009). They specify a DCC-MIDAS component model, in which short run fluctuations in correlations are captured by a standard DCC scheme and where the long run component is a weighted average of historical correlations. An advantage of this method is that it estimates long and short run correlations directly. However, it requires the estimation of weighting scheme parameters for both long run variances and long run correlations, which makes it harder to implement when the asset universe is large.

These high- and low-frequency volatilities are combined with a factor model to capture the high- and low-frequency dynamics of correlations. A distinct feature of the spline-GARCH model is that it allows unconditional volatility to vary over time, unlike traditional GARCH implementations in which it is constant. The factor model allows for a parsimonious representation of the correlation structure and makes it possible to distinguish between cycles in systematic volatility and patterns in idiosyncratic volatility. This is important because these fluctuations can have different asset pricing implications if they are driven by other underlying factors. The Factor-Spline-GARCH (FSG) model is a reduced form model designed to measure long and short run volatilities and correlations and as such it is not directly connected to stylized general equilibrium models such as those proposed by Bollerslev, Tauchen, and Zhou (2009) and Drechsler and Yaron (2011). However, Engle and Rangel (2008) show that the long-term component of volatility is related to macroeconomic fundamentals that affect future cash flows and discount rates, such as short-term interest rates, GDP growth, and inflation.

Following Rangel and Engle (2009), I specify a single-factor model to capture the covariance matrix of stock returns and to split returns into systematic and idiosyncratic components,

$$r_{it} = \alpha_i + \beta_i r_{Mt} + u_{it}, \quad (10)$$

The one-factor structure imposes a restriction on the covariance matrix by assuming that the idiosyncratic return u_{it} is uncorrelated with the market return and uncorrelated across stocks. However, when true betas are time-varying the assumption that the idiosyncratic return is uncorrelated with the market return no longer holds. Furthermore, it is likely that the single-factor structure is too simplistic and that latent factors affect returns and correlations. In that case the restriction that idiosyncratic returns are cross-sectionally uncorrelated is violated.

To capture these important features of the data and allow for richer correlation dy-

namics, Engle (2009b) allows for temporal deviations from these restrictions. In this extension, latent factors can influence conditional correlations and betas can vary over time but should be covariance-stationary and mean-revert to their unconditional expectation. This generalization yields the following expression for the conditional correlation

$$Cor_{t-1}(r_{it}, r_{jt}) = \frac{\beta_i \beta_j Var_{t-1}(r_{Mt}) + \beta_i E_{t-1}(r_{Mt} u_{jt}) + \beta_j E_{t-1}(r_{Mt} u_{it}) + E_{t-1}(u_{it} u_{jt})}{\sqrt{\beta_i^2 Var_{t-1}(r_{Mt}) + Var_{t-1}(u_{it}) + 2\beta_i E_{t-1}(r_{Mt} u_{it})} \sqrt{\beta_j^2 Var_{t-1}(r_{Mt}) + Var_{t-1}(u_{jt}) + 2\beta_j E_{t-1}(r_{Mt} u_{jt})}}. \quad (11)$$

The empirical implementation of (11) and decomposition of conditional correlations into high- and low-frequency terms requires estimates of long and short run market volatility and idiosyncratic volatility and estimates of correlations between the market return and idiosyncratic returns and among the idiosyncratic returns themselves.

As in Engle and Rangel (2008), I capture long and short run patterns in systematic volatility by writing market returns as

$$r_{Mt} = \alpha_M + u_{Mt} = \alpha_M + \sqrt{\tau_{Mt} g_{Mt}} \epsilon_{Mt} \quad \text{with} \quad \epsilon_{Mt} | \Phi_{t-1} \sim (0, 1), \quad (12)$$

where α_M is the unconditional market return, u_{Mt} the return innovation, and Φ_{t-1} the information set up to time $t-1$. g_{Mt} and τ_{Mt} are the high- and low-frequency components of market volatility, respectively. g_{Mt} reflects transitory volatility effects and τ_{Mt} captures long-term movements in volatility driven by changes in macroeconomic conditions.

The high-frequency component g_{Mt} is modeled as a unit GARCH process,

$$g_{Mt} = \left(1 - \theta_M - \phi_M - \frac{\gamma_M}{2}\right) + \theta_M \frac{(r_{Mt-1} - \alpha_M)^2}{\tau_{Mt-1}} + \gamma_M \frac{(r_{Mt-1} - \alpha_M)^2}{\tau_{Mt-1}} I_{r_{Mt-1} < 0} + \phi_M g_{Mt-1}, \quad (13)$$

where $I_{r_{Mt-1} < 0}$ is an indicator function for negative returns that captures the leverage effect in stock markets where negative returns have a stronger impact on volatility than

positive returns.

The low-frequency component τ_{Mt} is modeled as an exponential spline,

$$\tau_{Mt} = c_M \exp \left(w_{M0}t + \sum_{k=1}^{K_M} w_{Mk}((t - t_{Mk-1})_+)^2 \right), \quad (14)$$

where $(t - t_{Mk-1})_+ = (t - t_{Mk-1})$ if $t > t_{Mk-1}$ and 0 otherwise. K_M denotes the number of knots and determines the number of cycles in low-frequency volatility. Following Engle and Rangel (2008), I choose the optimal number of knots based on the Bayesian Information Criterion (BIC). The w_{Mk} coefficients determine the curvature of the volatility cycles.¹³

The specification of high- and low-frequency components of idiosyncratic volatility is similar to the decomposition of market volatility. In particular, I write the idiosyncratic return u_{it} as

$$\begin{aligned} u_{it} &= \sqrt{\tau_{it}g_{it}}\epsilon_{it} \quad \text{with} \quad \epsilon_{it}|\Phi_{t-1} \sim (0, 1) \\ g_{it} &= \left(1 - \theta_i - \phi_i - \frac{\gamma_i}{2}\right) + \theta_i \frac{(r_{it-1} - \alpha_i - \beta_i r_{Mt-1})^2}{\tau_{it-1}} + \gamma_i \frac{(r_{it-1} - \alpha_i - \beta_i r_{Mt-1})^2}{\tau_{it-1}} I_{r_{it-1} < 0} + \phi_i g_{it-1} \\ \tau_{it} &= c_i \exp \left(w_{i0}t + \sum_{k=1}^{K_i} w_{ik}((t - t_{ik-1})_+)^2 \right) \quad \forall i, \end{aligned} \quad (15)$$

Because the intensity of news and its effect on idiosyncratic risk depends on firm-specific conditions, I allow the number of cycles in long run idiosyncratic risk to vary across stocks based on the BIC.

I estimate the conditional correlations among the idiosyncratic returns and between these shocks and the market using the DCC approach of Engle (2002).¹⁴ Substituting these idiosyncratic correlations and the conditional expectations of high- and low-

¹³Because g_{Mt} is normalized to have an unconditional mean equal to one, unconditional volatility is equal to τ_{Mt} and time-varying in this model. The exponential form of the spline ensures that volatility is always positive.

¹⁴Engle (2009a) points out that when the factor model in equation (10) is correctly specified (i.e., beta is constant and there are no latent factors), all correlations produced by the DCC model should be zero.

frequency market volatility and idiosyncratic volatility into the expression for conditional correlations in equation (11), Rangel and Engle (2009) show that the high-frequency conditional correlation is given by

$$\begin{aligned} Cor_{t-1}(r_{it}, r_{jt}) = & \frac{\beta_i \beta_j \tau_{Mt} g_{Mt} + \beta_i \sqrt{\tau_{Mt} g_{Mt}} \sqrt{\tau_{jt} g_{jt}} \rho_{M,j,t}^\epsilon + \beta_j \sqrt{\tau_{Mt} g_{Mt}} \sqrt{\tau_{it} g_{it}} \rho_{M,i,t}^\epsilon + \sqrt{\tau_{it} g_{it}} \sqrt{\tau_{jt} g_{jt}} \rho_{i,j,t}^\epsilon}{\sqrt{\beta_i^2 \tau_{Mt} g_{Mt} + \tau_{it} g_{it} + 2\beta_i \sqrt{\tau_{Mt} g_{Mt} \tau_{it} g_{it}} \rho_{M,i,t}^\epsilon} \sqrt{\beta_j^2 \tau_{Mt} g_{Mt} + \tau_{jt} g_{jt} + 2\beta_j \sqrt{\tau_{Mt} g_{Mt} \tau_{jt} g_{jt}} \rho_{M,j,t}^\epsilon}}, \end{aligned} \quad (16)$$

where $\rho_{i,j,t}^\epsilon$ and $\rho_{M,i,t}^\epsilon$ are the correlation among idiosyncratic returns and the correlation between these shocks and the market, respectively.

Substitution of unconditional expectations from the expressions for volatilities and correlations into the unconditional version of (11) gives the low-frequency component of this correlation

$$\overline{Cor}_{t-1}(r_{it}, r_{jt}) = \frac{\beta_i \beta_j \tau_{Mt} + \sqrt{\tau_{it}} \sqrt{\tau_{jt}} \bar{\rho}_{i,j}^\epsilon}{\sqrt{\beta_i^2 \tau_{Mt} + \tau_{it}} \sqrt{\beta_j^2 \tau_{Mt} + \tau_{jt}}}. \quad (17)$$

Equation (17) shows that fluctuations in long run correlations reflect movements in long-term systematic and idiosyncratic variances. The high-frequency correlation in equation (16) mean-reverts to this time-varying low-frequency component. In contrast, in the standard DCC model conditional correlations mean-revert to constant unconditional correlations.

Finally, I combine the components of systematic and idiosyncratic volatility with the factor structure to construct high- and low-frequency components of total stock volatility

$$g_{it}^{total} = \beta_i^2 g_{Mt} + g_{it} \quad \text{and} \quad \tau_{it}^{total} = \beta_i^2 \tau_{Mt} + \tau_{it}. \quad (18)$$

4.3 Estimation of Volatility and Correlation Components

I estimate high- and low-frequency variances and correlations for the S&P 100 index and for the 30 components of the Dow Jones Industrial Average (DJIA) index as of February 18, 2008. The sample is daily and covers the period from January 1990 to December 2008. The FSG model is estimated using the two-step approach described by Rangel and Engle (2009) and parameter estimates are reported and discussed in Appendix A.

Figures 4 and 5 show the annualized high- and low-frequency components of market volatility and average idiosyncratic volatility, respectively. Both graphs reveal distinct cyclical patterns in volatility that highlight the importance of allowing unconditional volatility to vary over time. Consistent with Campbell, Lettau, Malkiel, and Xu (2001), I observe an upward trend in idiosyncratic risk until 2001. After 2001, however, idiosyncratic volatility decreases until 2006 and starts to rise again from 2007 onwards due to the credit crisis, which confirms results of Bekaert, Hodrick, and Zhang (2010).

[Figure 4 about here.]

[Figure 5 about here.]

The cycles in market volatility and average idiosyncratic volatility suggest that correlations also exhibit long-term trends. In particular, the one-factor structure in (10) implies that when systematic and idiosyncratic risk move in the same direction, they have opposite effects on correlations. An increase in market volatility leads to an increase in correlations whereas a rise in idiosyncratic risk results in a fall in correlations. In contrast, when aggregate and idiosyncratic volatility move in opposite directions, correlations can exhibit large movements. These patterns in market-wide correlations are illustrated in Figure 6. Because idiosyncratic risk rises until 1997 while systematic risk falls, average correlation levels drop substantially during this period. The credit crisis at the end of the sample period results in a sharp increase in average correlations, thereby reducing diversification opportunities when they are most needed.

[Figure 6 about here.]

5 Pricing of Long and Short Run Variance and Correlation Risk

5.1 Testing Methodology

The figures in the previous section show that market volatility, idiosyncratic risk, and correlations strongly fluctuate through time. Furthermore, the analysis reveals movements in volatilities and correlations at different frequencies. Changes in the high-frequency component of volatilities and correlations are transitory but movements in the low-frequency component are more permanent. In this section I examine whether these dynamics generate intertemporal hedging demands and significant risk premia in the cross-section of returns, as predicted by the ICAPM. I construct daily volatility and correlation factors as the first difference of the high- and low-frequency components of market volatility, average individual stock volatility, and average correlation obtained from the FSG model. I control for size, value, and momentum factors from the data library of Kenneth French.¹⁵ I also include a liquidity risk factor, which is calculated as the innovation in the value-weighted average of the Amihud (2002) illiquidity measure across all stocks in the CRSP universe.

Table 5 presents descriptive statistics for the daily volatility and correlation factors and for the other pricing factors. By construction, innovations in the high-frequency component of volatilities and correlations are strongly time-varying while shocks to the low-frequency part are more stable. Correlations between the long run and short run factors are low, which means that they can capture orthogonal sources of risk. The table further shows that innovations in high- (low-) frequency average stock volatility and average correlation are positively correlated with shocks to short (long) run market volatility, consistent with the decomposition of market volatility into individual volatilities and correlations. Pairwise correlations between the volatility and correlation factors and

¹⁵See http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html

the size, value, momentum, and liquidity factors are low. These factors can therefore have explanatory power for the cross-section of returns beyond the explanatory power of traditional risk factors.

[Table 5 about here.]

The ICAPM in equation (9) specifies a relation between returns and the *conditional* covariances of returns and risk factors. Bali (2008) and Bali and Engle (2010) point out that the use of time-varying conditional covariances is crucial in identifying a positive risk-return tradeoff. Furthermore, they stress that time series and cross-sectional data should be pooled to gain statistical power and find a significant risk-return relation. Pooling is possible because the ICAPM states that the intertemporal relation should be the same for all assets.

Following Bali and Engle (2010), I estimate time-varying conditional covariances between stock returns and risk factors using the standard DCC model. I pool time series and cross-sectional data by imposing the constraint that the γ and λ parameters in (9) are equal across stocks and estimate the resulting system of equations using the seemingly unrelated regression approach (SUR). The SUR model accounts for heteroskedasticity, autocorrelation, and contemporaneous correlation across residuals to estimate the system more efficiently.¹⁶

5.2 Estimates of the Price of Variance and Correlation Risk

I test whether long and short run market volatility risk, average individual volatility risk, and correlation risk are priced by examining the estimates of λ reported in Table 6. The

¹⁶A common concern in asset pricing tests is estimation error in factor exposures, i.e., the conditional covariances in equation (9). Since these covariances are obtained from a DCC model, standard methods to account for measurement error like the Shanken (1992) correction cannot be applied. Bali and Engle (2010) assess the impact of measurement error by comparing results of a one-step GARCH-in-mean estimation, in which the conditional covariances and prices of risk are estimated simultaneously, to results obtained from the two-step procedure that first estimates the covariances. They find that the one-step and two-step methods provide similar evidence on the risk-return relation and conclude that measurement error in conditional covariances does not have strong effects on the magnitude and statistical significance of the risk aversion coefficient.

first column corresponds to the traditional risk-return tradeoff without intertemporal hedging demand. The ICAPM implies that the estimated γ should be equal to the coefficient of relative risk aversion. The coefficient estimate is 2.86 with a t -statistic of 3.76, which is economically plausible and in line with the findings of Bali and Engle (2010). I also test the prediction of the ICAPM that all intercepts are equal to zero. The Wald statistic reported in the last row of the table shows that the hypothesis that all pricing errors are jointly zero cannot be rejected, which is evidence in favor of the conditional ICAPM.

Column two shows that the price of long run market volatility risk is significantly negative. The estimates in column three indicate that high-frequency market volatility risk is also negatively priced. Importantly, both factors remain significant when included in the regression together, consistent with the low correlation between the long- and short-term volatility factors reported in Table 5. The market factor also remains significant after including the two volatility factors. These findings extend the results of Adrian and Rosenberg (2008), who show that innovations in long- and short-term market volatility carry a negative price of risk in a cross-section of size-BE/ME portfolios, to a sample of individual stocks.

I study the pricing of innovations in high- and low-frequency components of average individual stock volatility in columns five and six. Both long and short run average volatility factors carry a significantly negative price of risk in the cross-section. Results are similar when the two average volatility factors enter the model simultaneously. Table 6 further shows that the price of the long and short run correlation risk factors is also negative and significant. Importantly, innovations in high- and low-frequency average volatilities and correlations continue to carry a negative price of risk when jointly added to the model. These findings suggest that aggregate volatility risk is priced because innovations to both average stock volatility and market-wide correlations are priced.

The negative prices of volatility and correlation risk imply that stocks that covary positively with unexpected changes in volatilities and correlations have a positive hedging

demand and therefore higher prices and lower expected returns. The negative prices of volatility and correlation risk in the cross-section of stocks are consistent with the positive variance and correlation risk premia in options reported in Table 1, because both reflect aversion of investors to sudden changes in volatilities and correlations.

The time series analysis in Section 3 further shows that the idiosyncratic part of variance premia in options is very different from the systematic component and does not have any predictive power for market returns. Therefore, I also decompose average stock volatility into systematic and idiosyncratic parts (see equation (18)) and add innovations in these components to the ICAPM. The results in Table 6 show that only the systematic component of average short-term volatility is significantly priced, in line with findings of Chen and Petkova (2010). However, the long-term systematic and idiosyncratic volatility risk factors are both priced.

[Table 6 about here.]

In Table 7 I study the pricing of long and short run volatility and correlation risk after controlling for size, value, momentum, and liquidity factors. The price of the *SMB*, *HML*, and *ILLIQ* factors has the expected sign but is insignificant, presumably because the 30 Dow stocks used as test assets are all large and liquid. The momentum factor is negatively priced but loses its significance after including the volatility risk factors. Both the long run and short run market volatility factor remains priced after controlling for the traditional set of risk factors. The innovations in average long and short run volatility and correlation also continue to carry significantly negative prices of risk that are similar to those reported in Table 6. The finding that the pricing of variance and correlation risk is unaffected by the inclusion of these other factors implies that the variance and correlation factors capture risks different from those related to the standard size, value, momentum, and liquidity factors.

[Table 7 about here.]

5.3 Factor Risk Premia

My empirical results show that market volatility risk is priced in the cross-section because innovations to average stock volatility and correlations are priced. In this section, I study the interaction between long and short run aggregate volatility risk, average volatility risk, and correlation risk in more detail. For each stock I calculate risk premia on these factors to assess their economic significance. These premia are computed by multiplying the factor exposures (conditional covariances) by the corresponding prices of market volatility risk in column four of Table 6 and the prices of average volatility risk and correlation risk in column 11. For ease of interpretation I aggregate the daily risk premia to a monthly frequency.

The results in Panel A of Table 8 show a wide spread in the premium for short run market volatility risk, ranging from -0.30% per month for Intel to 0.23% for Alcoa. In general, most growth stocks have a positive exposure to short-term market volatility risk, whereas most value stocks load negatively on this factor. Since the price of short run market volatility risk is negative, these loadings imply that the compensation for market volatility risk is negative for growth stocks and positive for value stocks. Thus, it appears that growth stocks earn lower returns than value stocks because they hedge against shocks to market volatility. Adrian and Rosenberg (2008) argue that the positive loadings of growth stocks on the short run market volatility factor could be driven by investor learning about the growth opportunities of these firms (see Pastor and Veronesi (2003)). Alternatively, the cross-sectional differences in exposures to volatility risk may reflect differences in the option value of growth opportunities due to heterogeneous adjustment costs (see Zhang (2005)). The premium for long run market volatility risk displays less cross-sectional variation but is sizeable for most stocks.

Column four shows that there is also a large spread in the compensation for short-term individual volatility risk. This is in line with the strong significance of the short run average volatility factor in the ICAPM regressions in Table 6. Similar to the premium for market volatility risk, the premium for average volatility risk is negative for most

growth stocks and positive for the majority of value stocks. The premium associated with long run average volatility risk is positive for all stocks. The monthly premium for short run correlation risk in column six ranges from -0.16 for Intel to 0.16 for Alcoa. The cross-sectional dispersion in the premium for long-term correlation risk is similar to the variation in the premium for long-term market volatility risk.

Panel B of Table 8 summarizes the relation between variance and correlation risk premia by showing their cross-sectional correlation. The premium for short-term market volatility risk is strongly positively correlated with the premia for short run average volatility risk and short run correlation risk. Thus, short-term market volatility risk is priced in the cross-section because both short-term individual volatility risk and short-term correlation risk are priced. The premium for long run market volatility risk is highly correlated with the premium for long run average volatility risk but also positively associated with the long run correlation risk factor.

Overall, the empirical results show that shocks to market volatility, average stock volatility, and market-wide correlations all carry significant negative prices of risk, consistent with the predictions of the ICAPM. Market volatility risk is priced because of both average volatility risk and correlation risk. Prices of risk are negative for shocks to both the long run and short run components of volatilities and correlations. The variance and correlation factors continue to have explanatory power after controlling for traditional size, value, momentum, and liquidity factors, which justifies their inclusion as separate factors in asset pricing models.

[Table 8 about here.]

5.4 Resolving the Idiosyncratic Volatility Puzzle

The analysis in the previous section focuses on the pricing of sudden changes in long- and short-term components of market volatility, average stock volatility, and correlations. In this section, I study the cross-sectional relation between stock returns and the level

of idiosyncratic risk motivated by the conflicting evidence on the pricing of idiosyncratic volatility presented by Ang, Hodrick, Xing, and Zhang (2006) and Fu (2009). I contribute to these papers by decomposing idiosyncratic risk into long- and short-term components using the Factor-Spline-GARCH model and adding these components to the ICAPM. Since my results show that volatility and correlation risk is priced, I control for the covariance of stock returns with these factors.¹⁷

Table 9 reports estimates of the price of long and short run components of idiosyncratic risk. The first column shows that the high-frequency, conditional measure of idiosyncratic volatility carries a significantly positive price of risk, consistent with Fu (2009). In contrast, the low-frequency, unconditional component of idiosyncratic volatility is negatively priced, which is in line with Ang, Hodrick, Xing, and Zhang (2006). Most importantly, both components continue to be priced when jointly included in the ICAPM. Thus, short-term idiosyncratic risk is priced differently from long-term idiosyncratic volatility.

After controlling for long and short run market volatility risk, high-frequency idiosyncratic volatility is no longer significantly priced. In contrast, the pricing of the low-frequency component is robust to the inclusion of the volatility risk factors. This result confirms the finding of Ang, Hodrick, Xing, and Zhang (2009) that exposure to long and short run market volatility risk does not explain the low returns on stocks with high unconditional idiosyncratic volatility. Including the long- and short-term average volatility risk factors leads to similar results. Consistent with Chen and Petkova (2010), I find that controlling for correlation risk does not explain the pricing of either long- or short-term idiosyncratic risk.

In sum, by decomposing idiosyncratic volatility into high- and low-frequency components I can reconcile the mixed evidence on the pricing of idiosyncratic risk. In particular, I find that the short-term component is positively priced while the long-term component carries a negative price of risk. The pricing of the conditional measure of idiosyncratic risk can be explained by exposure to volatility risk. In contrast, the significance of the

¹⁷Controlling for size, value, momentum, and liquidity factors does not affect the results.

price of unconditional idiosyncratic risk is robust to the inclusion of long and short run volatility and correlation risk factors.

[Table 9 about here.]

6 Conclusion

This paper examines the pricing of long and short run variance and correlation risk. At the aggregate level, I show that the predictive power of the market variance premium for stock market returns documented by Bollerslev, Tauchen, and Zhou (2009) is driven by the correlation risk premium and the systematic component of the variance premium in individual options. In contrast, idiosyncratic variance risk premia do contain information about future equity risk premia. Variance and correlation premia predict both monthly and quarterly returns and this predictive power is robust to the impact of the credit crisis and the inclusion of traditional return predictors.

I investigate the cross-sectional pricing of long and short run variance and correlation risk by decomposing market volatility, idiosyncratic volatility, and correlations into high- and low-frequency components using a novel approach proposed by Rangel and Engle (2009). I identify distinct cyclical patterns in long-term volatilities and correlations that highlight the importance of allowing unconditional volatilities and correlations to vary over time. The intertemporal CAPM predicts that these risk dynamics should be priced because investors want to hedge against a sudden increase in volatilities and correlations.

I find that innovations in long-term and short-term components of market volatility are negatively priced because shocks to average stock volatility and market-wide correlations are priced. Long and short run volatility and correlation risk factors capture different sources of risk and their cross-sectional explanatory power is not absorbed by size, value, momentum, and liquidity factors. Finally, I reconcile conflicting evidence on the pricing of idiosyncratic risk by showing that the long-term component of idiosyncratic volatility is negatively priced while the short-term component carries a positive price of risk.

Appendix A: Estimates for Spline-GARCH Model

In Table 10 I report parameter estimates for the Factor-Spline-GARCH model for the stocks in the Dow Jones and for the S&P 100 index.¹⁸ All ARCH coefficients in column one are significant at the 5% level except those for Intel and Coca-Cola, which are significant at the 10% level. Individual ARCH effects range from 0.01 to 0.13 with a mean of 0.06. In contrast, the ARCH effect for the market index is close to zero and insignificant. GARCH estimates are significant for all stocks and the market factor at a 1% level. The GARCH effect is between 0.22 and 0.98 for individual stocks with a mean value of 0.77 and equals 0.90 for the market index. As pointed out by Rangel and Engle (2009), the large cross-sectional variation in ARCH and GARCH effects indicates that there is substantial variation in the persistence of idiosyncratic volatilities across stocks.

Estimates of the leverage effects are in column three. The leverage coefficient is significant for about half of the stocks and ranges from 0.00 to 0.12 with a mean of 0.05. The leverage effect for the market index is much stronger than for most individual stocks. Column four reports the estimated degrees of freedom of the Student's t -distribution, which are between 4 and 9 with a mean of 6. These values indicate excess kurtosis, which highlights the importance of assuming a t -distribution for return innovations in the estimation. The optimal number of knots in the spline function varies from 1 to 9 for individual stocks and is equal to 6 for the market index. Since the sample covers 19 years and the knots are equally spaced, this means that the length of each market volatility cycle is just over 3 years. The large variation in the number of knots across firms reveals important differences in the cyclical patterns in their idiosyncratic volatilities.

Panel B in Table 10 presents estimates of the second-stage DCC parameters, which are both significant at a 1% level. α_{DCC} equals 0.0028 and β_{DCC} is 0.9920, indicating strong persistence in the correlation between idiosyncratic returns.

[Table 10 about here.]

¹⁸I do not report alphas and betas. Alphas are small and only significant at a 5% level for AIG and GM. Market betas are centered around one and range from 0.62 for Johnson & Johnson to 1.39 for Intel.

References

- Adrian, T., and J. Rosenberg, 2008, Stock returns and volatility: Pricing the long-run and short-run components of market risk, *Journal of Finance* 63, 2997–3030.
- Amihud, Y., 2002, Illiquidity and stock returns: Cross-section and time-series effects, *Journal of Financial Markets* 5, 31–56.
- Andersen, T., T. Bollerslev, F. Diebold, and H. Ebens, 2001, The distribution of realized stock return volatility, *Journal of Financial Economics* 61, 43–76.
- Ang, A., R. Hodrick, Y. Xing, and X. Zhang, 2006, The cross-section of volatility and expected returns, *Journal of Finance* 61, 259–299.
- Ang, A., R. Hodrick, Y. Xing, and X. Zhang, 2009, High idiosyncratic volatility and low returns: International and further U.S. evidence, *Journal of Financial Economics* 91, 1–23.
- Bali, T., 2008, The intertemporal relation between expected return and risk, *Journal of Financial Economics* 87, 101–131.
- Bali, T., and R. Engle, 2010, The intertemporal capital asset pricing model with dynamic conditional correlations, *Journal of Monetary Economics* 57, 377–390.
- Bansal, R., and A. Yaron, 2004, Risks for the long run: A potential resolution of asset pricing puzzles, *Journal of Finance* 59, 1481–1509.
- Bekaert, G., and E. Engstrom, 2010, Asset return dynamics under bad environment-good environment fundamentals, Working Paper, Columbia University.
- Bekaert, G., R. J. Hodrick, and X. Zhang, 2010, Aggregate idiosyncratic volatility, Working Paper, Columbia University.
- Black, F., and M. Scholes, 1973, The pricing of options and corporate liabilities, *Journal of Political Economy* 81, 637–654.

- Bollerslev, T., R. Engle, and D. Nelson, 1994, ARCH models, in R. Engle, and D. McFadden, ed.: *Handbook of Econometrics* . pp. 2959–3038 (North-Holland: Amsterdam).
- Bollerslev, T., R. Engle, and J. Wooldridge, 1988, A capital asset pricing model with time-varying covariances, *Journal of Political Economy* 96, 116–131.
- Bollerslev, T., M. Gibson, and H. Zhou, 2009, Dynamic estimation of volatility risk premia and investor risk aversion from option-implied and realized volatilities, *Journal of Econometrics*, Forthcoming.
- Bollerslev, T., G. Tauchen, and H. Zhou, 2009, Expected stock returns and variance risk premia, *Review of Financial Studies* 22, 4463–4492.
- Britten-Jones, M., and A. Neuberger, 2000, Option prices, implied price processes, and stochastic volatility, *Journal of Finance* 55, 839–866.
- Campbell, J., and J. Cochrane, 1999, By force of habit: A consumption based explanation of aggregate stock market behavior, *Journal of Political Economy* 107, 205–251.
- Campbell, J., M. Lettau, B. Malkiel, and Y. Xu, 2001, Have individual stocks become more volatile? An empirical exploration of idiosyncratic risk, *Journal of Finance* 56, 1–43.
- Carr, P., and L. Wu, 2009, Variance risk premiums, *Review of Financial Studies* 22, 1311–1341.
- Chen, Z., and R. Petkova, 2010, Does idiosyncratic volatility proxy for risk exposure, Working Paper, Texas A&M University.
- Chernov, M., R. Gallant, E. Ghysels, and G. Tauchen, 2003, Alternative models for stock price dynamics, *Journal of Econometrics* 116, 225–257.
- Colacito, R., R. Engle, and E. Ghysels, 2009, A component model for dynamic correlations, Working Paper, University of North Carolina at Chapel Hill.

- Drechsler, I., and A. Yaron, 2011, What's vol got to do with it, *Review of Financial Studies* 24, 1–45.
- Driessen, J., P. Maenhout, and G. Vilkov, 2009, The price of correlation risk: Evidence from equity options, *Journal of Finance* 64, 1377–1406.
- Engle, R., 2002, Dynamic conditional correlation: A simple class of multivariate generalized autoregressive conditional heteroskedasticity models, *Journal of Business and Economic Statistics* 20, 339–350.
- Engle, R., 2009a, *Anticipating Correlations: A New Paradigm for Risk Management* (Princeton University Press: Princeton).
- Engle, R., 2009b, High dimension dynamic correlations, in J. Castle, and N. Shephard, ed.: *The Methodology and Practice of Econometrics: A Festschrift in Honour of David F. Hendry* . pp. 122–148 (Oxford University Press: Oxford).
- Engle, R., and G. Lee, 1999, A long run and short run component model of stock return volatility, in R. Engle, and H. White, ed.: *Cointegration, Causality, and Forecasting: A Festschrift in Honour of Clive W. J. Granger* . pp. 475–497 (Oxford University Press: Oxford).
- Engle, R., and J. Rangel, 2008, The spline-GARCH model for low-frequency volatility and its global macroeconomic causes, *Review of Financial Studies* 21, 1187–1222.
- Fu, F., 2009, Idiosyncratic risk and the cross-section of expected stock returns, *Journal of Financial Economics* 91, 24–37.
- Ghysels, T., A. Harvey, and E. Renault, 1996, Stochastic volatility, in C. Rao, and G. Maddala, ed.: *Handbook of Statistics* . pp. 119–191 (North-Holland: Amsterdam).
- Goyal, A., and P. Santa-Clara, 2003, Idiosyncratic risk matters!, *Journal of Finance* 58, 975–1007.

- Guo, H., and R. Savickas, 2008, Average idiosyncratic volatility in G7 countries, *Review of Financial Studies* 21, 1259–1296.
- Hodrick, R., 1992, Dividend yields and expected stock returns: Application to financial markets, *Review of Financial Studies* 5, 357–386.
- Jiang, G., and Y. Tian, 2005, The model-free implied volatility and its information content, *Review of Financial Studies* 18, 1305–1342.
- Krishnan, C., R. Petkova, and P. Ritchken, 2009, Correlation risk, *Journal of Empirical Finance* 16, 353–367.
- Lettau, M., and S. Ludvigson, 2001, Consumption, aggregate wealth, and expected stock returns, *Journal of Finance* 56, 815–849.
- Merton, R., 1973, An intertemporal asset pricing model, *Econometrica* 41, 867–887.
- Merton, R., 1987, A simple model of capital market equilibrium with incomplete information, *Journal of Finance* 42, 483–510.
- Moskowitz, T., 2003, An analysis of covariance risk and pricing anomalies, *Review of Financial Studies* 16, 417–457.
- Pastor, L., and P. Veronesi, 2003, Stock valuation and learning about profitability, *Journal of Finance* 58, 1749–1789.
- Pollet, J., and M. Wilson, 2010, Average correlation and stock market returns, *Journal of Financial Economics* 96, 364–380.
- Rangel, J., and R. Engle, 2009, The factor-spline-GARCH model for high and low frequency correlations, Working Paper, New York University.
- Schurhoff, N., and A. Ziegler, 2010, The pricing of systematic and idiosyncratic variance risk, Working paper, University of Lausanne.

- Shanken, J., 1992, On the estimation of beta-pricing models, *Review of Financial Studies* 5, 1–34.
- Todorov, V., 2009, Variance risk premium dynamics: The role of jumps, *Review of Financial Studies*, Forthcoming.
- Zhang, L., 2005, The value premium, *Journal of Finance* 60, 67–103.

Table 1: Summary Statistics for Variances and Correlations

This table reports descriptive statistics for monthly stock market returns, market variances, individual variances, and correlations for the period from January 1996 to December 2008. R_M is the log return on the S&P 100 in excess of the three-month T-bill rate. IV_M is the model-free implied market variance computed from index options. RV_M is the realized market variance calculated by summing squared five-minute returns on the S&P 100 index over a one-month window. VRP_M is the market variance risk premium, defined as the difference between IV_M and RV_M . $\overline{IV}$ is the value-weighted average model-free implied variance for individual stocks computed from stock options. $\overline{RV}$ is the average variance risk premium on individual options, defined as the difference calculated by summing squared five-minute stock returns. $\overline{VRP}$ is the implied market variance and the implied variance of the stocks in the S&P 100 index. RC is the realized correlation, calculated from the realized market variance and the realized variance of the stocks in the S&P 100 index. CRP is the correlation risk premium, defined as the difference between IC and RC . Panel A reports the mean, median, standard deviation, skewness, kurtosis and autocorrelation for these variables and Panel B presents their pairwise correlations.

	R_M	IV_M	RV_M	VRP_M	$\overline{IV}$	$\overline{RV}$	$\overline{VRP}$	IC	RC	CRP
<i>Panel A: Summary Statistics</i>										
Mean	0.10	37.90	26.50	11.40	101.18	117.14	-15.95	0.36	0.21	0.15
Median	0.66	32.81	15.99	10.24	77.44	87.84	0.32	0.35	0.20	0.15
Std. Dev.	4.71	33.59	44.15	24.77	71.99	116.14	63.33	0.11	0.08	0.09
Skewness	-0.69	3.21	6.91	-4.74	1.85	4.09	-6.55	0.27	0.72	0.42
Kurtosis	3.91	17.42	63.30	51.60	7.63	29.29	59.56	2.68	3.73	2.99
AR(1)	0.06	0.68	0.55	0.19	0.76	0.63	0.33	0.81	0.62	0.80
<i>Panel B: Correlation Matrix</i>										
R_M	1.00									
IV_M	-0.45	1.00								
RV_M	-0.39	0.83	1.00							
VRP_M	0.08	-0.12	-0.66	1.00						
$\overline{IV}$	-0.30	0.80	0.76	-0.26	1.00					
$\overline{RV}$	-0.33	0.80	0.92	-0.56	0.88	1.00				
$\overline{VRP}$	0.28	-0.55	-0.82	0.72	-0.47	-0.84	1.00			
IC	-0.08	0.52	0.34	0.09	0.14	0.17	-0.16	1.00		
RC	-0.37	0.58	0.54	-0.18	0.36	0.34	-0.21	0.66	1.00	
CRP	0.24	0.15	-0.05	0.29	-0.14	-0.08	-0.01	0.72	-0.05	1.00

Table 2: **Predicting Monthly Returns with Variance and Correlation Risk Premia**

This table presents results for predictability regressions for the period from January 1996 to December 2008. The dependent variable is the monthly excess return on the S&P 100 index in percent. In Panel A the independent variables are lagged VRP_M , $\overline{VRP}$, and CRP , all defined in Table 1. Also included are $\overline{VRP}^{syst}$, which is the value-weighted average of the systematic part of the variance risk premium on individual options, and $\overline{VRP}^{idio}$, which is the value-weighted average of the idiosyncratic variance risk premium. Panel B adds traditional return predictors to the regressions. CAY is the consumption-wealth ratio, for which the most recently available quarterly observations are taken as monthly observations, DEF is the default spread, defined as the yield differential between bonds rated Baa by Moody's and bonds with a Moody's rating of Aaa, $\log(P/E)$ is the log of the smoothed price-earnings ratio for the S&P 500, and $TERM$ is the term spread, defined as the difference between the ten-year and three-month Treasury yield. All independent variables are standardized to have mean zero and standard deviation equal to one. Intercept estimates are not reported to save space. Robust t -statistics based on Hodrick (1992) standard errors are in parentheses.

<i>Panel A: Predictability Regressions Excluding Traditional Predictors</i>								
VRP_{Mt-1}	0.94 (5.12)				0.07 (0.11)			
$\overline{VRP}_{t-1}$		0.85 (2.63)		0.87 (2.93)	0.81 (1.09)			
CRP_{t-1}			0.94 (2.67)	0.95 (2.66)	0.93 (2.29)			
$\overline{VRP}_{t-1}^{syst}$						0.90 (4.47)		0.80 (2.67)
$\overline{VRP}_{t-1}^{idio}$							0.56 (0.99)	0.29 (0.64)
Adj. R^2 (%)	3.35	2.62	3.36	6.13	5.52	3.04	0.88	2.76
<i>Panel B: Predictability Regressions Including Traditional Predictors</i>								
VRP_{Mt-1}	0.75 (3.01)				0.02 (0.04)			
$\overline{VRP}_{t-1}$		0.80 (2.62)		0.77 (1.85)	0.70 (1.11)			
CRP_{t-1}			1.03 (2.04)	0.99 (1.98)	0.97 (1.86)			
$\overline{VRP}_{t-1}^{syst}$						0.74 (2.89)		0.59 (1.29)
$\overline{VRP}_{t-1}^{idio}$							0.62 (1.15)	0.38 (0.68)
CAY_{t-1}	0.41 (1.47)	0.68 (2.37)	-0.11 (-0.25)	0.09 (0.18)	0.09 (0.19)	0.40 (1.46)	0.72 (2.35)	0.57 (1.69)
DEF_{t-1}	-0.69 (-1.72)	-0.57 (-1.57)	-0.70 (-1.87)	-0.30 (-0.54)	-0.32 (-0.57)	-0.72 (-1.80)	-0.69 (-1.76)	-0.58 (-1.57)
$\log(P/E)_{t-1}$	-0.31 (-0.81)	-0.17 (-0.49)	-0.00 (-0.01)	-0.07 (-0.16)	-0.01 (-0.02)	-0.35 (-0.89)	-0.00 (-0.01)	-0.23 (-0.48)
$TERM_{t-1}$	-0.07 (-0.14)	-0.20 (-0.46)	-0.38 (-0.92)	-0.49 (-1.02)	-0.52 (-1.00)	-0.05 (-0.10)	-0.26 (-0.56)	-0.17 (-1.00)
Adj. R^2 (%)	3.64	3.93	4.11	5.60	4.67	3.53	2.92	3.32

Table 3: **Predicting Quarterly Returns with Variance and Correlation Risk Premia**

This table presents results for predictability regressions for the period from January 1996 to December 2008. The dependent variable is the quarterly excess return on the S&P 100 index in percent. The quarterly regressions are based on overlapping monthly returns. All independent variables are as defined in Table 2 and are standardized to have mean zero and standard deviation equal to one. Intercept estimates are not reported to save space. Robust t -statistics based on Hodrick (1992) standard errors are in parentheses.

<i>Panel A: Predictability Regressions Excluding Traditional Predictors</i>								
VRP_{Mt-1}	2.61 (4.73)					1.69 (1.18)		
$\overline{VRP}_{t-1}$		1.58 (2.34)		1.63 (2.73)		0.38 (0.26)		
CRP_{t-1}			2.66 (3.20)	2.69 (3.24)		2.17 (2.42)		
$\overline{VRP}_{t-1}^{syst}$						2.45 (4.98)		2.54 (3.60)
$\overline{VRP}_{t-1}^{idio}$							0.71 (0.48)	-0.24 (-0.23)
Adj. R^2 (%)	8.55	2.72	8.94	11.93	12.82	7.47	0.01	6.93
<i>Panel B: Predictability Regressions Including Traditional Predictors</i>								
VRP_{Mt-1}	1.37 (2.03)					1.17 (0.70)		
$\overline{VRP}_{t-1}$		0.46 (0.85)		0.36 (0.75)		-0.49 (-0.37)		
CRP_{t-1}			2.80 (2.44)	2.85 (2.46)		2.42 (1.95)		
$\overline{VRP}_{t-1}^{syst}$						1.24 (1.75)		1.48 (1.91)
$\overline{VRP}_{t-1}^{idio}$							-0.04 (-0.03)	-0.62 (-0.79)
CAY_{t-1}	0.98 (1.50)	1.24 (1.87)	-0.51 (-0.59)	-0.52 (-0.58)	-0.53 (-0.60)	0.99 (1.50)	1.11 (1.60)	0.72 (1.00)
DEF_{t-1}	-3.71 (-4.19)	-4.04 (-4.43)	-3.47 (-4.20)	-3.24 (-3.76)	-3.35 (-3.97)	-3.80 (-3.69)	-4.31 (-4.79)	-4.02 (-4.34)
$\log(P/E)_{t-1}$	-0.27 (-0.38)	0.06 (0.06)	0.39 (0.41)	0.40 (0.40)	0.08 (0.06)	-0.30 (-0.28)	0.09 (0.06)	-0.49 (-0.64)
$TERM_{t-1}$	0.08 (0.10)	-0.00 (-0.01)	-0.77 (-0.68)	-0.82 (-0.74)	0.08 (0.06)	0.11 (0.09)	0.08 (0.06)	0.31 (0.38)
Adj. R^2 (%)	19.31	17.51	22.84	22.47	22.54	18.93	17.29	18.72

Table 4: **Pre-Crisis Return Prediction with Variance and Correlation Risk Premia**

This table presents results for predictability regressions for the pre-crisis period from January 1996 to December 2007. The dependent variable in Panel A is the monthly excess return on the S&P 100 index in percent. Panel B shows forecasting results for quarterly regressions that are based on overlapping monthly returns. All independent variables are as defined in Table 2 and are standardized to have mean zero and standard deviation equal to one. Intercept estimates are not reported to save space. Robust t -Statistics based on Hodrick (1992) standard errors are in parentheses.

<i>Panel A: Monthly Return Regressions</i>								
VRP_{Mt-1}	0.92 (2.33)					0.44 (1.03)		
$\overline{VRP}_{t-1}$		-0.07 (-0.08)		0.05 (0.07)		-0.14 (-0.17)		
CRP_{t-1}			0.86 (2.56)	0.86 (2.54)		(0.72) (2.02)		
$\overline{VRP}_{t-1}^{syst}$						0.74 (1.92)		0.66 (1.40)
$\overline{VRP}_{t-1}^{idio}$							-0.47 (-0.78)	-0.27 (-0.44)
Adj. R^2 (%)	1.18	-0.02	3.23	2.54	2.14	0.73	-0.02	0.19
<i>Panel B: Quarterly Return Regressions</i>								
VRP_{Mt-1}	3.29 (3.73)					2.46 (0.70)		
$\overline{VRP}_{t-1}$		-0.93 (-0.35)		-0.57 (-0.44)		-1.64 (-1.09)		
CRP_{t-1}			2.46 (2.51)	2.44 (2.54)		1.69 (2.02)		
$\overline{VRP}_{t-1}^{syst}$						2.48 (2.56)		2.12 (2.51)
$\overline{VRP}_{t-1}^{idio}$							-1.92 (-1.16)	-1.26 (-1.02)
Adj. R^2 (%)	7.56	-0.03	10.47	9.97	12.55	4.91	2.37	5.45

Table 5: **Summary Statistics for Risk Factors**

This table presents descriptive statistics for daily variance and correlation risk factors and for traditional pricing factors. R_M is the excess return on the S&P 100. $\Delta HF MVOL$ and $\Delta LF MVOL$ are innovations in the high- and low-frequency components of market volatility, respectively. $\Delta HF AVOL$ and $\Delta LF AVOL$ are innovations in the value-weighted average of high- and low-frequency volatility of all stocks in the Dow Jones index. $\Delta HF COR$ and $\Delta LF COR$ are innovations in the weighted average of high- and low-frequency correlations between the stocks in the Dow Jones. Innovations in high- and low-frequency components of volatilities and correlations are obtained from the Factor-Spline-GARCH model described in Section 4.2 and are multiplied by 100 for ease of exposition. SMB , HML , and UMD are the daily size, value, and momentum factors from the data library of Kenneth French. $\Delta ILLIQ$ is the innovation in the value-weighted average of the Amihud (2002) illiquidity measure across all stocks in the CRSP universe. The sample is daily and covers the period from January 1990 to December 2008.

	Mean	Std. Dev.	R_M	ΔHF MVOL	ΔLF MVOL	ΔHF AVOL	ΔLF AVOL	ΔHF COR	ΔLF COR	SMB	HML	UMD	$\Delta ILLIQ$
R_M	0.0128	1.17	1.00										
$\Delta HF MVOL$	0.0215	8.81	0.09	1.00									
$\Delta LF MVOL$	0.0127	0.11	-0.02	0.01	1.00								
$\Delta HF AVOL$	0.0378	6.89	0.07	0.88	0.02	1.00							
$\Delta LF AVOL$	0.0375	0.22	-0.02	0.01	0.92	0.02	1.00						
$\Delta HF COR$	0.0004	2.69	0.05	0.88	0.00	-0.05	0.00	1.00					
$\Delta LF COR$	-0.0020	0.02	0.00	0.00	0.37	0.00	-0.11	0.01	1.00				
SMB	0.0036	0.57	-0.25	-0.08	-0.01	-0.06	-0.01	-0.10	0.01	1.00			
HML	0.0161	0.58	-0.33	0.00	-0.01	-0.01	-0.02	0.01	0.02	-0.16	1.00		
UMD	0.0458	0.82	-0.20	0.00	0.02	-0.01	0.01	0.00	0.00	0.09	-0.17	1.00	
$\Delta ILLIQ$	-0.0010	0.49	-0.09	-0.07	0.00	-0.07	0.00	-0.03	0.00	-0.01	0.01	0.00	1.00

Table 6: Cross-Sectional Price of Long and Short Run Volatility and Correlation Risk

This table reports estimates of the cross-sectional price of long and short run volatility and correlation risk. The intertemporal risk-return tradeoff is estimated as a system of equations for a panel of 30 Dow Jones stocks using the SUR approach. The dependent variable is the daily excess stock return and the independent variables are the conditional covariances between stock returns and risk factors. The factors are the daily excess market return and the high- and low-frequency volatility and correlation factors defined in Table 5. The average volatility factor is further decomposed into systematic and idiosyncratic parts. The coefficients for the high- and low-frequency correlation factors have been divided by 10 for ease of presentation. The sample period is from January 1990 to December 2008. *t*-Statistics adjusted for heteroskedasticity, autocorrelation, and cross-correlation among the residuals are in parentheses. *Wald* is the Wald test statistic for the null hypothesis that all intercepts are jointly equal to zero.

R_M	2.86 (3.76)	2.30 (2.83)	3.58 (4.58)	3.05 (3.66)	3.55 (4.57)	2.40 (3.05)	3.16 (3.90)	3.11 (4.07)	3.09 (4.02)	3.31 (4.28)	3.09 (3.80)	3.05 (3.94)	2.43 (3.00)
Δ HF MVOL		-1.08 (-4.08)		-1.06 (-4.01)									
Δ LF MVOL			-89.60 (-2.00)	-82.72 (-1.85)									
Δ HF AVOL					-0.91 (-4.19)	-0.86 (-3.96)					-0.77 (-3.37)		
Δ LF AVOL						-33.24 (-2.21)	-26.11 (-1.72)				-42.29 (-2.29)		
Δ HF COR								-0.60 (-1.95)		-0.53 (-1.72)	-0.45 (-1.69)		
Δ LF COR									-125.41 (-2.67)	-123.21 (-2.58)	-150.38 (-4.38)		
Δ HF AVOL _{syst}												-1.08 (-3.92)	
Δ LF AVOL _{syst}												-87.29 (-1.98)	
Δ HF AVOL _{idio}													0.70 (0.93)
Δ LF AVOL _{idio}													-30.25 (-1.75)
Wald	15.70 (0.99)	18.42 (0.95)	17.00 (0.97)	19.70 (0.99)	17.35 (0.97)	16.99 (0.97)	18.27 (0.95)	18.29 (0.95)	22.70 (0.83)	24.07 (0.77)	30.24 (0.45)	19.66 (0.93)	16.92 (0.97)
<i>p</i> -value													

Table 7: **Price of Volatility and Correlation Risk Controlling for Traditional Factors**

This table reports estimates of the cross-sectional price of long and short run volatility and correlation risk. The intertemporal risk-return tradeoff is estimated as a system of equations for a panel of 30 Dow Jones stocks using the SUR approach. The dependent variable is the daily excess stock return and the independent variables are the conditional covariances between stock returns and risk factors. The factors are the daily excess market return, the high- and low-frequency volatility and correlation factors, and the size, value, momentum, and illiquidity factors defined in Table 5. The average volatility factor is further decomposed into systematic and idiosyncratic parts. The coefficients for the high- and low-frequency correlation factors have been divided by 10 for ease of presentation. The sample period is from January 1990 to December 2008. *t*-Statistics adjusted for heteroskedasticity, autocorrelation, and cross-correlation among the residuals are in parentheses. *Wald* is the Wald test statistic for the null hypothesis that all intercepts are jointly equal to zero.

R_M	2.03 (2.08)	2.14 (2.15)	1.80 (1.74)	2.14 (2.20)	2.59 (2.61)
Δ HF MVOL		-1.05 (-3.92)			
Δ LF MVOL		-82.20 (-1.99)			
Δ HF AVOL			-0.83 (-3.81)		-0.74 (-3.21)
Δ LF AVOL			-29.83 (-1.82)		-38.76 (-2.19)
Δ HF COR				-0.53 (-1.72)	-0.51 (-1.70)
Δ LF COR				-132.95 (-2.86)	-150.01 (-3.26)
SMB	2.08 (0.84)	2.38 (1.00)	3.94 (0.83)	1.30 (0.54)	2.19 (-0.46)
HML	1.57 (0.67)	1.20 (0.50)	1.07 (0.45)	3.04 (1.28)	2.29 (0.96)
MOM	-4.29 (-2.00)	-3.06 (-1.33)	-3.13 (-1.37)	-4.30 (-1.99)	-2.19 (-0.94)
ILLIQ	-5.66 (-0.86)	-6.96 (-1.06)	-5.45 (-0.83)	-4.44 (-0.67)	-2.77 (-0.42)
Wald	20.10	22.23	20.40	29.10	31.03
p-value	(0.91)	(0.87)	(0.91)	(0.51)	(0.41)

Table 8: **Factor Risk Premia**

This table reports risk premia for all stocks in the Dow Jones index on the market factor and the volatility and correlation factors. The premia are computed by multiplying factor exposures by the corresponding prices of market risk and market volatility risk (both from column four in Table 6) and the prices of average volatility risk and correlation risk (from column 11 in Table 6). The daily risk premia are aggregated to a monthly frequency for ease of interpretation.

	R_M	Δ HF MVOL	Δ LF MVOL	Δ HF AVOL	Δ LF AVOL	Δ HF COR	Δ LF COR
<i>Panel A: Factor Risk Premia</i>							
AA	0.48	0.23	0.16	0.13	0.09	0.16	0.03
AIG	0.59	0.20	0.30	0.12	0.24	0.11	0.13
AXP	0.69	0.14	0.15	0.07	0.08	0.07	0.02
BA	0.38	0.08	0.21	0.01	0.16	0.06	0.11
BAC	0.58	0.12	0.24	0.01	0.17	0.09	0.09
C	0.80	0.17	0.23	0.02	0.13	0.13	0.11
CAT	0.46	0.17	0.21	0.04	0.13	0.13	0.06
CVX	0.22	0.01	0.15	0.02	0.07	-0.02	0.07
DD	0.42	0.09	0.16	0.02	0.07	0.06	0.06
DIS	0.47	0.04	0.17	-0.04	0.08	0.02	0.09
GE	0.56	-0.07	0.15	-0.09	0.07	-0.03	0.07
GM	0.56	0.13	0.21	0.00	0.13	0.10	0.06
HD	0.62	0.10	0.19	0.02	0.07	0.06	0.20
HPQ	0.63	-0.04	0.10	-0.11	0.03	0.00	0.00
IBM	0.42	-0.11	0.08	-0.11	0.03	-0.08	0.00
INTC	0.81	-0.30	0.13	-0.23	0.07	-0.16	0.07
JNJ	0.22	0.00	0.15	0.01	0.07	-0.02	0.08
JPM	0.74	0.11	0.13	0.01	0.02	0.08	0.02
KO	0.26	-0.06	0.16	-0.04	0.08	-0.04	0.11
MCD	0.26	0.13	0.11	0.07	0.03	0.09	0.06
MMM	0.28	0.13	0.18	0.04	0.12	0.07	0.08
MRK	0.35	-0.03	0.12	-0.04	0.02	-0.02	0.04
MSFT	0.61	-0.10	0.15	-0.11	0.05	-0.05	0.15
PFE	0.39	0.14	0.15	0.09	0.07	0.07	0.10
PG	0.24	-0.09	0.17	-0.09	0.09	-0.06	0.07
T	0.35	0.05	0.12	-0.01	0.05	0.01	0.03
UTX	0.40	0.22	0.15	0.08	0.10	0.15	0.03
VZ	0.32	-0.06	0.11	-0.11	0.04	-0.03	0.03
WMT	0.45	-0.08	0.11	-0.11	0.01	-0.04	0.11
XOM	0.27	-0.12	0.15	-0.07	0.07	-0.09	0.07

Panel B: Cross-Sectional Correlation of Factor Risk Premia

R_M	1.00						
Δ HF MVOL	0.07	1.00					
Δ LF MVOL	0.25	0.52	1.00				
Δ HF AVOL	-0.12	0.94	0.47	1.00			
Δ LF AVOL	0.13	0.49	0.94	0.46	1.00		
Δ HF COR	0.16	0.98	0.52 ₄₅	0.87	0.49	1.00	
Δ LF COR	0.09	0.03	0.54	0.08	0.34	0.02	1.00

Table 9: **Cross-Sectional Pricing of Long and Short Run Idiosyncratic Risk**

This table reports estimates of the cross-sectional price of long and short run idiosyncratic risk. The intertemporal risk-return tradeoff is estimated as a system of equations for a panel of 30 Dow Jones stocks using the SUR approach. The dependent variable is the daily excess stock return and the independent variables are the high- and low-frequency components of idiosyncratic volatility. I control for the exposures to the market factor and the volatility and correlation risk factors defined in Table 5. The coefficients for the high- and low-frequency correlation factors have been divided by 10 for ease of presentation. The sample period is from January 1990 to December 2008. *t*-Statistics adjusted for heteroskedasticity, autocorrelation, and cross-correlation among the residuals are in parentheses.

R_M	3.00 (3.80)	3.46 (4.46)	3.14 (3.98)	3.36 (4.00)	3.52 (4.26)	3.53 (4.43)	3.49 (4.22)
Δ HF MVOL				-1.08 (-4.07)			
Δ LF MVOL				-113.33 (-2.36)			
Δ HF AVOL					-0.87 (-3.97)		-0.79 (-3.44)
Δ LF AVOL					-34.80 (-2.19)		-63.25 (-3.78)
Δ HF COR						-0.51 (-1.68)	-0.14 (-0.44)
Δ LF COR						-123.01 (-4.57)	-156.01 (-5.52)
HF IVOL	0.02 (2.01)		0.03 (2.15)	0.02 (1.36)	0.02 (1.35)	0.03 (2.30)	0.02 (1.10)
LF IVOL		-0.06 (-3.97)	-0.08 (-4.47)	-0.08 (-4.63)	-0.08 (-4.50)	-0.08 (-4.46)	-0.08 (-4.54)

Table 10: **Spline-GARCH-DCC Estimates**

This table reports estimation results for the Spline-GARCH model in Panel A and the DCC model in Panel B for all stocks in the Dow Jones and for the S&P 100 index. For stock i , θ_i is the ARCH effect, ϕ_i the GARCH effect, γ_i the leverage effect, ν_i the degrees of freedom of the Student's t -distribution, and K_i the optimal number of knots in the spline function. α_{DCC} and β_{DCC} are the DCC parameters. The sample is daily and covers the period January 1990 to December 2008.

	(1)		(2)		(3)		(4)	(5)
	θ_i	t -stat	ϕ_i	t -stat	γ_i	t -stat	ν_i	K_i
<i>Panel A: Spline-GARCH Estimates</i>								
AA	0.02	3.47	0.94	62.17	0.02	2.03	6	4
AIG	0.06	4.08	0.82	30.26	0.07	3.61	7	7
AXP	0.05	3.66	0.83	23.87	0.06	3.25	6	6
BA	0.04	2.36	0.84	10.78	0.06	2.54	6	4
BAC	0.12	5.40	0.65	13.08	0.08	2.66	6	8
C	0.09	4.31	0.74	15.98	0.09	3.46	6	4
CAT	0.08	3.98	0.40	2.94	0.00	0.04	5	9
CVX	0.05	4.90	0.91	55.57	0.02	1.16	9	3
DD	0.07	2.94	0.75	6.65	0.00	0.03	6	5
DIS	0.10	3.89	0.56	6.25	0.04	1.25	5	4
GE	0.04	2.61	0.84	15.30	0.06	3.05	7	4
GM	0.02	2.53	0.96	107.59	0.03	4.52	5	4
HD	0.02	2.74	0.91	32.71	0.05	3.64	6	4
HPQ	0.01	3.85	0.98	214.39	0.00	0.24	4	4
IBM	0.07	3.32	0.68	11.38	0.12	3.52	4	6
INTC	0.02	1.90	0.93	42.28	0.03	2.61	5	4
JNJ	0.04	3.56	0.83	29.27	0.12	4.85	6	4
JPM	0.03	2.51	0.88	23.41	0.07	3.84	6	5
KO	0.04	1.75	0.81	17.51	0.10	3.35	6	4
MCD	0.03	4.84	0.96	165.41	0.01	0.64	6	1
MMM	0.13	4.57	0.22	2.44	0.00	0.01	5	8
MRK	0.09	3.26	0.44	4.56	0.09	2.18	5	8
MSFT	0.08	4.36	0.77	19.45	0.06	2.19	5	4
PFE	0.08	5.63	0.81	22.41	0.01	0.91	6	5
PG	0.05	3.09	0.79	14.18	0.06	2.52	6	8
T	0.07	4.71	0.84	26.51	0.01	0.87	7	4
UTX	0.08	3.55	0.60	5.84	0.08	2.35	6	6
VZ	0.11	4.50	0.67	7.72	0.01	0.83	7	4
WMT	0.03	5.40	0.96	150.16	0.00	0.52	6	1
XOM	0.08	5.24	0.86	35.36	0.00	0.18	8	3
SP100	0.00	0.61	0.90	70.90	0.13	8.31	9	6
<hr/>								
	α_{DCC}	t -stat	β_{DCC}	t -stat				
<hr/>								
<i>Panel B: DCC Estimates</i>								
DCC	0.0028	16.51	0.99	1331.41				
<hr/>								

Figure 1: **Market Variance Risk Premium**

This figure plots the implied and realized variance (top panel) and the variance risk premium (bottom panel) for the S&P 100 index from January 1996 to December 2008. The shaded areas indicate NBER recession periods.

Figure 2: **Average Variance Risk Premium**

This figure plots the value-weighted average of the implied and realized variances (top panel) and of the variance risk premia (bottom panel) for the stocks in the S&P 100 index from January 1996 to December 2008. The average variance risk premium is further decomposed into a systematic and an idiosyncratic part. The shaded areas indicate NBER recession periods.

Figure 3: **Correlation Risk Premium**

This figure plots the average implied and realized correlation between the stocks in the S&P 100 index (top panel) and the correlation risk premium (bottom panel) from January 1996 to December 2008. The shaded areas indicate NBER recession periods.

Figure 4: Long and Short Run Market Volatility

This figure plots the high- and low-frequency components of market volatility for the period January 1990 to December 2008. The high-frequency component is modeled as a unit GARCH process and the low-frequency component is modeled as an exponential quadratic spline as described in Section 4.2. The estimation is based on daily returns on the S&P 100 index and the resulting volatilities are annualized.

Figure 5: **Average Long and Short Run Idiosyncratic Volatility**

This figure plots the value-weighted average of the high- and low-frequency components of idiosyncratic volatility of the stocks in the Dow Jones index for the period January 1990 to December 2008. The high-frequency component is modeled as a unit GARCH process and the low-frequency component is modeled as an exponential quadratic spline as described in Section 4.2. The estimation is based on daily stock returns and the resulting volatilities are annualized.

Figure 6: **Average Long and Short Run Correlation**

This figure plots the weighted average of the high- and low-frequency components of the correlations between all stocks in the Dow Jones index for the period January 1990 to December 2008. The high- and low-frequency correlations are obtained from the Factor-Spline-GARCH model described in Section 4.2.