Asymmetric Volatility Risk:
Evidence from Option Markets[*]

Jens Jackwerth   **Grigory Vilkov**

This version: September 23, 2013

Abstract

Using non-parametric methods to model the dependencies between risk-neutral distributions of the market index (S&P 500) and its expected volatility (VIX), we show how to extract the expected risk-neutral correlation between the index and its future expected volatility. Comparing the implied correlation to its realized counterpart reveals a significant risk premium priced into the index-volatility correlation, which can be interpreted as the compensation for the fear of rising volatility during and after a market crash, i.e., fear of crash continuation for a prolonged period of time. We show how the index-volatility correlation premium is related to future market returns and explain its economics.

Keywords: Asymmetric volatility, SPX options, VIX options, implied correlation

JEL: G11, G12, G13, G17

1 Introduction

As the finance profession moved beyond the simplistic Black-Scholes (1973) model, the importance of volatility has been widely accepted, e.g., in the GARCH models, e.g. Bauwens, Laurent, and Rombouts (2006), and in stochastic volatility models starting with Heston (1993). Surprisingly little is however known about the bivariate distribution of returns and volatility. GARCH models view postulate a deterministic dependence, and stochastic volatility models work with the assumptions of a fixed instantaneous correlation, which is then assumed to be the same under the risk-neutral and the actual distribution. That latter assumption then leaves no room for a price of correlation risk between returns and volatility, contrary to our empirical findings. We contribute to the understanding of the dependence between returns and volatilities by working out the risk-neutral bivariate distribution from option prices on the index and on volatility. Song and Xiu (2013) deemed this feat impossible as there are no cross-asset options written on both market index and its volatility, but we use information from longer dated index options to achieve just that. Figure 1 explains our approach in a simplified risk-neutral setting. The first period return distribution (1.3 and 0.7 with 50% probability each) is given by the index options. The second period volatility distribution ($\sigma_{high} = 0.4$ and $\sigma_{low} = 0.1$) is given by the options on volatility. We now vary the dependence from perfect positive correlation (Panel A) to perfect negative correlation (Panel B). It emerges that the two-period distribution is affected by the dependence structure (namely, the skewness changes sign), and we use market prices of long-dated index options to calibrate this dependence.

[Figure 1 about here]

In practice, we use one- and two-month maturity options on the S&P 500 and one-month options on VIX plus the relevant futures and interest rates data. We assume that the second period return distributions inherit the shape of the first period distribution but for its volatility that is drawn from the option implied volatility distribution. The dependence structure is modeled by a Frank copula, which is particularly well-suited for the typically negative dependence

we find in the data. We obtain in the end an option-implied measure of index-to-volatility correlation as well as complete information about the bivariate distributions for each month from May 2006 to January 2013. We also compute implied index-to-volatility correlation under the actual distribution and further the correlation premium as the difference between risk-neutral and actual values; this quantity will typically be negative.

Our findings contribute in four main ways to our understanding of index returns and volatilities. First, we add to our understanding of the interaction between returns and volatilities. The index-to-volatility correlation is always negative and varies considerably over time. More negative implied index-to-volatility correlations lead to more left-skewed two-month returns distributions (more precisely, they lead to a fatter left tail in the second-month distribution compared to the first-month one), and hence to more crash-o-phobic investor attitudes and a steeper two-month implied volatility smile. Thus, negative market returns translate into a prolonged fear of a continuing crash. A less negative implied correlation suggests less left-skewed two-month distributions and a lower fear of a continuing crash. Alternatively, we express this mechanism as steepening of the pricing kernel (projected onto the index returns) from the first month to the second month. Such steepening also goes hand-in-hand with more pessimistic investor attitudes about market down turns.

Second, the index-to-volatility correlation premium is statistically significant and predicts future index returns. This predictive power is a new finding, and it is an additional predictive component, not to be confounded with the variance risk premium. The predictive power is about half of the variance risk premium and survives in joint predictive regressions. Interestingly, while our implied index-to-volatility correlation measure picks up investor attitudes towards left-tail events, adding option-implied skewness does not improve the predictive regressions.

Third, we provide information on the factors driving the index-to-volatility correlation premium, namely autocorrelation in VIX. High autocorrelation in VIX leads to a significantly more negative index-to-volatility correlation premium, which means a more pessimistic investor outlook onto the future. Volatility (option implied or historical), volatility of historical variance,

and forward-looking implied correlation among the equities of the S&P 500 do not explain the index-to-volatility correlation premium by themselves, but when autocorrelation in VIX and an interaction term are added, results turn strongly significant. The economics of this finding are that assessments of risk need to be combined with autocorrelation in order to explain investor attitudes to future down turns, as measured by the implied index-to-volatility correlation premium. In particular, the high risk measures lead to a more negative index-to-volatility correlation premium, the same holds for the autocorrelation of VIX. However, the interacted terms are significant and positive, thus making the index-to-volatility correlation premium less negative. This new finding suggests that joint negative effects on investor attitudes of high risk and high autocorrelation are less than the sum of the individual effects. The interaction term then compensates for this difference. We thus find some mean reversion in the combined pessimist outlook due to high risk and high VIX autocorrelation: each fact dims the investors view of the future but taken together, the future outlook reverts somewhat back towards a more optimistic view.

Fourth, knowledge about the bivariate distribution of returns and volatilities matters for risk management and option pricing. In particular the time-varying nature of the finding suggests a time-varying and priced role for index-to-volatility correlation risk.

Section 2 relates the paper to the literature. All data are presented in Section 3. The paper develops the methodology in Section 4. Hypotheses and results follow in Section 5. Section 6 concludes. Technical details are are left for appendix.

2 Literature

We investigate in this paper the relationship of returns and volatility, and we back out information thereabout from option prices and historical prices. This relates us directly to two large areas of research. First, the asymmetric volatility area on the interaction of returns and volatility and, second, the recovery of market price implied information.

2.1 Asymmetric volatility and return predictability

A current, state-of-the-art account of asymmetric volatility is Bekaert and Wu (2000) with the main modeling choice that returns and volatility are inverse related to each other, as empirically found in the data. This informs for example on trading of volatility, which can serve as a hedge to return exposure. Whaley (2013) is a nice account of trading in volatility products. Closely related is the literature on the variance risk premium, which is defined as the risk-neutral expectation of index variance minus the actual index variance. It measures fear of crashes, of high volatility, and of tail risks. Select contributions to this field are Bekaert and Hoerova (2012), Bollerslev, Tauchen, and Zhou (2009), Bollerslev and Todorov (2011), and Carr and Wu (2009). Some of these papers then go on to predicting future returns using the variance risk premium within an ICAPM framework. Alternative factors used in this setting are market skewness by Chang, Christoffersen, and Jacobs (2012), jump (crash) risk by Cremers, Halling, and Weinbaum (2012), and rare disasters by Gabaix (2012).

2.2 Recovery of market price implied information

We are interested in recovering the implied correlation between returns and volatility. This is not to be confused with the implied correlation among stocks within the index. We do not relate to that literature other than that we share the word correlation.

We touch on recovering marginal densities of S&P 500 returns and of future VIX values, separately. While we go on to analyzing the joint distribution, there is a large literature on the extraction of marginal densities, see the surveys of Jackwerth (2004) or Christoffersen, Jacobs, and Chang (2012) for the risk-neutral aspect and any number of time-series investigations on estimating actual distributions. Song and Xiu (2013) infer risk-neutral S&P 500 and VIX risk-neutral distributions, separately, while stating that the bivariate distribution cannot be obtained; a point that we see differently. Allen, McAleer, Powell, Singh (2012) estimate actual S&P 500 and VIX distributions, separately. Again, no attempt is made to estimate the bivariate

distribution. Chabi-Yo and Song (2013) go one step further and try to combine the two marginal distributions with the historical correlation coefficient. They argue this shortcut by observing that in continuous time models, the risk-neutral and the actual correlation coefficients are instantaneously the same. But the whole point of estimating the bivariate distributions over longer horizons is that this identity then no longer holds and the index-to-volatility correlation premium will not be zero any longer.

A number of papers approach the estimation of the risk-neutral correlation between returns and volatility by resorting to parametric, continuous time models with constant instantaneous correlation coefficients. Bardgett, Gourier, and Leippold (2013), Duan and Yeh (2012), Branger and Voelkert (2013), and Carr and Madan (2013) work along those lines. Fuertes and Papanicolaou (2012) is an econometric approach to estimating the Heston (1993) model. A common feature of all these models is the tightly circumscribed parametric model which is estimated. Also, these models fix the instantaneous correlation while we look at the correlation of one-month returns with one-month expected volatility.

There exist a parallel literature which estimates the correlations based on the actual returns and volatilities, often drawing on high-frequency data, e.g. Ishida, McAleer, and Oya (2011). Zhang, Mykland, and Ait-Sahalia (2005) is a high-frequency estimation of volatilities but, at the suggestion of one of the authors, can also be used to estimating correlations from high-frequency time series data.

Shen (2013) does not deal outright with the bivariate distribution of returns and volatility but estimates two volatility term structures. One he uncovers from the SPX index options and the other from futures on VIX volatilities. He documents empirical inconsistencies between the two volatility term structures.

3   Data

During our sample period from January 2006 to January 2013, we collect information on futures and options on the S&P 500 and on VIX.

3.1   S&P 500 data

We obtain end-of-day options data on S&P 500 from OptionMetrics. We also obtain tick data for S&P 500 futures from TickData. We work with midpoint implied volatilities inferred from raw option prices. We also record the S&P 500 dividend yields and interpolate the zero certificate of deposit rates from OptionMetrics to match the exact days-to-maturity of our S&P 500 options.

We find the day closest to 25 days before the expiration date of a particular month. This typically turns out to be a date between the 20th and the 28th day of the previous months; this filter gives us the most uniform time-to-maturity for the nearest and the second nearest month. We eliminate options with zero bids and filter for moneyness (=strike price/index level) to lie between 0.7 and 1.3. We use only out-of-the-money options. Standing on this observation date, we collect the one-month and two-month index options. We also compute model-free implied variance from one-month and model-free implied skewness from one- and two-month options using the results of Bakshi, Kapadia, and Madan (2003).

3.2   VIX data

We obtain end-of-day options data on VIX from OptionMetrics. We also obtain daily VIX futures data from CBOE. The underlying VIX we collect as daily data from CBOE and as intraday data from TickData. We work with midpoint implied volatilities. We estimate VIX implied volatilities from the Black (1976) model using raw option prices and reported VIX

futures level at the end of the day, as the implied volatilities reported by OptionMetrics are incorrectly based on the actual VIX level without adjusting for the cost of carry.[1]

We use the same observation date as for the S&P 500 options. We eliminate zero bids and filter for moneyness (=strike price/VIX futures level) to lie between 0.2 and 2.2; note that in the case of VIX we go very far out-of-the-money, and the reason is that VIX has much higher implied volatility than index options, and the options are liquid farther from the ATM level. We use only out-of-the-money options. Standing on this observation date, we collect the one-month VIX options.

3.3 Realized index-to-volatility correlation

We use the second-best method (resampling and averaging) by Zhang, Mykland, and Ait-Sahalia (2005) to compute realized index-to-volatility correlations from intraday data on S&P 500 futures and VIX levels. Using high-frequency data improves estimates (see e.g. Andersen, Bollerslev, Diebold, and Ebens (2001), and Barndorff-Nielsen and Shephard (2002)) but microstructure effects can render these estimates inconsistent. One the one hand, noise effects arising from bid-ask bounce or discrete trading drive realized variances to infinity, while, on the other hand, the effect of nonsynchronicity, known as the Epps (1979) effect, drives covariances to zero. Reducing the sampling frequency can mitigate both problems. To avoid wasting any data and still having a good estimator we follow a modified version of the second-best method of Zhang, Mykland, and Ait-Sahalia (2005): we subsample and average first at frequencies of 5, 10, 20, 30, and 60 minutes and then average these five estimators. For each observation date on which we estimate the risk-neutral densities for S&P 500 and VIX, we compute the realized index-to-volatility correlation using the data over the past month (30 calendar days). Following Bollerslev, Tauchen, and Zhou (2009), we also compute the variance risk premium for S&P 500 as a difference between currently observed model-free implied variance and realized variance over the past month computed from the high frequency data.

---

[1]We used the IvyDB OptionMetrics database available through WRDS, and updated on March 11, 2013 (as noted on WRDS web-site) until 01/2013.

4 Methodology

We turn to estimating the marginal risk-neutral distributions for the index and the volatility next. Thereafter we discuss the estimation of the joint distributions before we turn to our hypotheses.

4.1 Marginal distribution for S&P 500

To estimate the marginal S&P 500 one-month distribution, we use fast and stable method by Jackwerth (2004), which provides (given a tradeoff parameter) a closed form solution for fitting the implied volatilities of observed options best while at the same time delivering the smoothest implied volatility smile. The same paper argues that, given some low number of observed option prices, the exact choice of method matters little in order to obtain risk-neutral distributions. We require at least eight index option prices to exist at any given observation date. We obtain implied volatilities over the moneyness (=strike price/index level) range from 0.7 to 1.3 and interpolate/extrapolate them to fill in a fine grid within 0.5 to 1.5 moneyness range.[2] We improve the fit of the risk-neutral distribution by forcing the mean to be the risk-free rate minus the expected dividend yield (obtained from OptionMetrics) and the volatility such that the observed option prices are matched best in a root mean squared error sense. Measuring the root mean squared error of implied volatilities does not make much difference.

4.2 Marginal distribution for VIX

We also use the method of Jackwerth (2004) to obtain the risk-neutral distributions of one-month volatility options for dates where we have at least 5 available options. We obtain implied volatilities for all available out-of-the-money (i.e., with moneyness = strike price/futures level above one for calls and below one for puts) options, and we interpolate and extrapolate the implied volatilities to fill in a fine grid within the 0.2 to 2.2 moneyness range. We do not force

---

[2]For robustness, we changed the moneyness range from 0.7-1.3 to 0.8-1.2 and results do not change by much. Also, for increased numerical stability, we multiply all moneyness levels by 1000 before applying the method and retranslating afterwards.

the mean as to a particular value, as that value is not predetermined since volatility is not a traded asset. We improve the fit of the risk-neutral distribution by forcing the volatility such that the observed option prices are matched best in a root mean squared error sense (and again, measuring the error for implied volatilities does not change the results). However, both effects are very small compared to the noise in the VIX options data and hardly affect our fit at all.

[Figure 2 about here]

The original option prices, the resulting risk-neutral densities, and the comparison of market prices and model prices can be studied in Figure 2 for one particular observation date in our sample.

4.3   Joint distribution of returns and volatilities

After having the two marginal distributions for returns and future expected volatilities in place, we model the bivariate distribution via a Frank (1979) copula.[3] This copula works best for our situation of typically sizeable negative correlation between the two variables. The Frank copula has a single parameter, which controls the correlation between the two variables.

A further degree of freedom lies in the choice of the mean of the risk-neutral volatility distribution. Since the underlying volatility is not a traded asset, the mean of the distribution is not determined by the cost of carry and thus not equal to the value of the one-month future on VIX. We tackle this problem by simply estimating the mean of the risk-neutral volatility distribution and the parameter of the Frank copula from the fit of the two-month index option prices. Technical details on the procedure are relegated to an appendix.

In order to obtain the two-month index option prices, we proceed as follows. We pick starting values for the parameter of the Frank copula and for the mean of the risk-neutral volatility distribution. That allows us to draw 10,000 bivariate draws from the Frank copula where

---

[3]See Patton (2009) on copulas in finance and Trivedi and Zimmer (2007) for a detailed guide which includes the Frank copula in particular.

9

each draw consists of two uniformly distributed values. These we feed into the interpolated cumulative distribution functions of the risk-neutral returns distribution and the risk-neutral volatility distribution, where we reset the mean according to our starting value. We now have 10,000 draws of a one-month return and of a second month volatility (i.e. the volatility that we expect to prevail over the second month).

Next, we construct the second period return distribution for each joint draw of one-month return and second-month volatility. Here, we make the assumption that the second period return distribution keeps the shape of the first period return distribution but with changed mean and volatility. The second period mean is simply the second period risk-free rate minus the expected dividend yield. The second period volatility we obtain from each draw from the Frank copula. The first period return distribution has been estimated as a discrete distribution on a fine grid of return values. We demean the returns and divide by the standard deviation in order to normalize the distribution, before multiplying by the second period volatility and adding back the second period mean. Thus, for each draw, we now have a discrete distribution of second period returns. Multiplication of each of these returns with the first period return gives a distribution of two period returns with probability equal to the probability of drawing jointly the first period index return and the respective second period expected volatility times the probability of the second period return from the discrete distribution.

Repeating the above procedure for each of our draws from the Frank copula, we arrive at a complete two period risk-neutral return distribution. Our first-period index distribution has on average 250 realizations, and hence we simulate about 10,000 times 250 = 2,500,000 two-period returns.[4] We can now obtain two-month option prices, one for each observed two-month out-of-the-money option. As a loss measure, we use the root-mean squared error of the implied volatilities of model prices and observed prices. We finally calibrate our starting parameter for

---

[4]In order to reduce computational demands, we map the two-period marginal distribution of the index onto a reduced set of 500 values each by aggregating neighboring probabilities.

the Frank copula and our second period mean of the volatility distribution such that we are minimizing our loss measure.[5]

[Figure 3 about here]

[Figure 4 about here]

We depict the bivariate distribution of one-month S&P500 returns and volatilities in Figure 3 for the same observation date which we use in Figure 2 for the marginal distributions. We set the parameter of the Frank copula such that we obtain a correlation of -0.87 between returns and volatilities. The parameter and the mean of the second period volatility have been chosen so that the root mean squared error of the implied volatilities for the two-month index options (model vs. observed within a 0.8 to 1.2 moneyness range) has been minimized. For the same observation date, we show in Figure 4 the two-month index option prices. We depict the observed bid and ask prices, as well as the model prices which lie within the bid and ask prices.

5   Hypotheses and Results

Our findings contribute in three main ways to our understanding of index returns and volatilities.

5.1   Interaction between returns and volatilities

First, we add to our understanding of the interaction between returns and volatilities. The correlation between index returns and volatilities is always negative and varies considerably over time. As we can see in Figure 5, this holds for the risk-neutral as well as the actual index-to-volatility correlations. The correlation between the two time-series of correlations is 0.0259.

---

[5]Alternative choices of root mean squared error or mean absolute error combined with relative or absolute measures combined with implied volatilities or prices yield similar results.

[Figure 5 about here]

The mean of the risk-neutral correlations is -0.8137 and the mean of the actual correlations is -0.7361. This finding allows us a first hypothesis:

H1: There exists a risk-premium related to the index-to-volatility correlation.

A t-test shows a strongly significant negative risk premium of -0.0776 with a t-statistic of -3.6744. For an economic interpretation, this means that holding index-to-volatility correlation risk is being rewarded in the market. We would like to think about the index-to-volatility correlation risk in terms of a stochastically changing investment opportunity set and argue that our correlation risk measure tells us something about future crash fears and their duration.

We thus investigate which exact risk it is that is being compensated here. The point of departure of our analysis is the variance risk premium, which is the payment for a hedge against negative returns by gaining on increased volatility. Inherently, the variance risk premium is symmetric and a large up and a large down move (relative to the mean) result in the same contribution to variance.[6] A first indication that the index-to-volatility correlation risk premium is measuring something different lies in the low correlation with the variance risk premium of only 0.03. Our economic intuition is that the index-to-volatility correlation risk premium rather works asymmetrically in that a more negative risk premium worsens the investment opportunity set by strengthening the relationship between a negative first period shock and a more volatile second period return which leads to a more left-skewed second period distribution which in turns leads to a more left-skewed two-period distribution. We are now ready to formulate our second hypothesis related to tail fatness:

H2: An increase in the ratio of left tail probability beyond a certain threshold for the second
period over the first period (tail fatness) and a decrease in the slope of the ratio of first period

---

[6]Bollerslev and Todorov (2011) claim that the variance risk premium is mostly driven by jumps and not purely by the diffusive part of the returns. However, their measures of jumps for the left tail and the right tail are correlated to a very high degree. Our argument that the variance risk premium is about symmetric tail risks thus remains valid.

over second period pricing kernels (shift in pricing kernels) should be negatively related to the

index-to-volatility correlation risk premium.

We first test our hypothesis by creating our measure of tail fatness ($tail - fatness_t$) based on the ratio of left tail probability beyond the 10th percentile return (-7.77% monthly in our sample) for the second period over the first period.[7] We then run the following regression, including one or more of the regressors:

$$tail - fatness_t = \alpha + \beta_{VRP}VRP_t + \beta_{CRP}CRP_t + \beta_{SKEW}SKEW_t + \varepsilon_t, \qquad (1)$$

where VRP_t is the variance risk premium, CRP_t is the index-to-volatility correlation premium, and $SKEW_t$ is the first period implied skewness computed as in Bakshi, Kapadia, and Madan (2003). Results can be found in Table 1, Panel A.

[Table 1 about here]

We find confirmation for our second hypothesis in model (1) of Table 1, Panel A, where we regress tail fatness solely on a constant and the index-to-volatility correlation premium. The coefficients have zero p-values and the adjusted R-squared is almost 16%. Using only the variance risk premium with a constant in model (2) gives an insignificant coefficient for the variance risk premium and a low adjusted R-squared of 2%. More interesting is model (3) where we use all variables and add the risk-neutral skewness of the first period return distribution. The index-to-volatility correlation premium remains important (p-value of zero) and the variance risk premium is only marginally significant (p-value of 11%). However, implied skewness has an important negative influence (p-value of zero) on tail fatness. For interpretation recall that implied skewness for the index is typically negative, and hence a more left-skewed distribution in the first month, i.e., where the implied skewness is more negative, is associated with a less pronounced index-to-volatility correlation premium. Effectively, the tail fears of investors are already incorporated in the first-month distribution, and we do not expect the situation to become even worse next period.

---

[7]We base our threshold return on the 10th percentile of first period returns. Results are qualitatively similar when we use the 5th or 25th percentile instead.

A more negative implied index-to-volatility correlation premium thus leads to a fatter left tail in the second-month distribution compared to the first-month distribution. This will increase fears of a crash-o-phobic investor about second period crashes, which in turn leads typically to a steeper two-month implied volatility smile. Negative market returns then translate into a prolonged fear of a continuing crash.

Second, we define the slope of the ratio of first period over second period pricing kernels (shift in pricing kernels). For that we compute the first and second month risk-neutral cumulative probability densities at 20 return points corresponding to the specific percentile of the first-month return distribution (we use percentiles 5 to 100 with increment of 5). Then we approximate the probability density function (i.e., dQ) as the probability mass assigned to the returns within each bin defined by the stipulated percentiles (i.e., returns between percentiles 0 and 5, 5 and 10, and so on). Assuming that the actual distribution does not change from one month to the next, we compute the ratio of the pricing kernels as the ratio of the second months to the first months risk-neutral distribution, depicted in Figure 6.

[Figure 6 about here]

We compute the slope between the probability ratio in 0 to 5th percentile bin and the 25th to 30th percentile bin to be our $shift_t$ variable. This variable moves exactly opposite to tail fatness: if tail fatness is high, then slope is lower (more negative). We then run the following regression, including one or more of the regressors:

$$shift_t = \alpha + \beta_{VRP}VRP_t + \beta_{CRP}CRP_t + \beta_{SKEW}SKEW_t + \varepsilon_t, \tag{2}$$

where VRP_t is the variance risk premium, CRP_t is the index-to-volatility correlation premium, and $SKEW_t$ is the first period implied skewness computed as in Bakshi, Kapadia, and Madan (2003). Results for estimation of the regression (2) can be found in Table 1, Panel B. The results are much the same as the ones for tail fatness, except that all sign are flipped. A more negative implied index-to-volatility correlation premium thus leads a steeper pricing kernel in the second-month distribution compared to the first-month distribution and thus to a more

14

negative slope ($shift_t$). This will again increase fears of a crash-o-phobic investor about second period crashes which in turn leads typically to a steeper two-month implied volatility smile. Negative market returns then translate into a prolonged fear of a continuing crash.

5.2 Prediction of future index returns

The index-to-volatility risk premium seems to measure time-varying investment opportunities according to our above analysis. Thus, we are curious to investigate if it also predicts future index returns. Namely, a worsening of the future investment opportunities (more negative index-to-volatility risk premium) should lead to a higher required rate of return in order to compensate the investor for the dimmer outlook.

Again, we start our analysis with the variance risk premium, which is known to predict index returns, see e.g. Bollerslev, Tauchen, and Zhou (2009). Also, since we found implied skewness to matter for tail fatness, but the mechanism was very mechanical. In particular, implied skewness seemed to measure the current shape of the distribution more than the future investment opportunity set. We are thus skeptical if implied skewness predicts future index returns. We thus formulate our hypothesis.

H3: Future S&P 500 index returns are being predicted by the variance risk premium and the index-to-volatility risk premium. Implied skewness does not predict future S&P 500 index returns

We test our hypothesis by running the following regression:

$$r_{S\&P500,t+1} = \alpha + \beta_{VRP}VRP_t + \beta_{CRP}CRP_t + \beta_{SKEW}SKEW_t + \varepsilon_t, \qquad (3)$$

where $r_{S\&P500,t+1}$ are the future 25-day returns on the S&P 500 from observation date until the maturity of the index options. VRP_t is the variance risk premium, CRP_t is the index-to-volatility correlation premium, and $SKEW_t$ is the first period implied skewness as computed by Bakshi, Kapadia, and Madan (2003). Results of the estimation of regression (3) can be found in Table 2.

In model (1), we find that our index-to-volatility correlation risk premium predicts future index returns. The p-value of the coefficient is 0.0176 and the adjusted R-squared is almost 4%. We further confirm in model (2) that the variance risk premium predicts future index returns. The p-value of the coefficient is 0.0098 and the adjusted R-squared is just above 8%.

The variance risk premium and the index-to-volatility correlation risk premium both predict returns, and are almost uncorrelated with a coefficient of only 0.03 in our sample. To confirm that the two variables account for different information, we run the joint model (3) where we also add implied skewness. Both premiums survive strongly significantly, while implied skewness has a p-value of 0.5560. The R-squared of the joint regression is almost 12% (some 0.72% higher if we drop implied skewness from the regression), indicating that the joint model almost adds the explanatory power of models (1) and (2).

5.3   Drivers of the index-to-volatility correlation premium

The index-to-volatility correlation premium expresses market expectation of the future and hence should to some extent be based on past market information. If a more negative index-to-volatility correlation premium indicates worsening future investment opportunities (e.g., higher risk or duration of down markets), then we might capture this effect also by historical variables such as autocorrelation in VIX, historical volatility, VIX, the volatility of historical variance, or option implied correlation among the equities of the S&P 500. We compute autocorrelation in VIX as the sum of the first three lags in the regression of daily VIX values on its lagged values over the past quarter; historical S&P 500 return volatility is derived as the square root of the realized variance over the past month computed from high-frequency data, volatility of historical variance is computed from daily variances of S&P 500 returns over the past week, where each daily variance is again computed from intraday data using the resampling and

16

averaging algorithm discussed in the data section. Implied correlations among the equities of the S&P 500 index are estimated following Driessen, Maenhout, and Vilkov (2013).

[Table 3 about here]

There is a significant negative dependency between the autocorrelation in VIX and the index-to-volatility correlation premium in Model (1) of Table 3. Thus, persistence in VIX leads to worsening future investment opportunities (a more negative index-to-volatility correlation premium). This finding is intuitive for high levels of VIX, which suggest persistence in volatile times which is clearly a bad perspective. Less clear is the story for low VIX levels. Thus, we turn to our risk measures and investigate them in turn.

Historical volatility ($PastSP500Volatility_t$) is insignificant by itself in explaining the index-to-volatility correlation premium; see Table 3, Model (2). But as we are concerned that the full story about future investment opportunities only emerges once VIX autocorrelation is also taken into account, we next investigate the following interacted regression:

$$CRP_t = \alpha + \beta_{AC}AC_t + \beta_{PastSP500Volatility}PastSP500Volatility_t+$$

$$\beta_{interaction}PastSP500Volatility_t AC_t + \varepsilon_t, \qquad (4)$$

where CRP_t is the implied index-to-volatility risk premium. AC_t is the VIX autocorrelation and $PastSP500Volatility_t$ is historical volatility on the index. Results for (4) are in Model (3) of Table 3 and show that VIX autocorrelation keeps its significant negative effect. Past volatility also has now a significant negative effect, suggesting that a risky environment leads to worsening future investment opportunities. However, the interacted term is significant and positive, thus making the index-to-volatility correlation premium less negative. This novel finding suggests that the combined negative effects on future investment opportunities of high risk and high autocorrelation are less than the sum of the individual effects. Effectively, investors have faced bad times characterized by high volatility and steadily high fear index VIX for some time, and they might expect a slowly decreasing risk over the next months, i.e., some kind of a mean-reversion that is typical for such risk measures as volatility.

The same story holds for the remaining risk measures, namely VIX (Models (4 and 5)), volatility of variance (Models (6 and 7)), and the option implied correlation among the equities constituting the S&P 500 index (Models (8 and 9)).

5.4 Bivariate distribution of returns and volatilities

As we have estimates of the bivariate distributions of returns and volatilities, we can use the risk-neutral distributions right away for pricing of exotic options written on a combination of the two quantities. Such knowledge could also inform the risk management of banks holding assets linked to either quantity about the risk neutral correlation between the two quantities. Furthermore, the time-varying nature of the finding suggests a time-varying and priced role for index-to-volatility correlation risk. This is relevant for the development of stochastic volatility option pricing models, which mostly model the correlation as constant.

6 Conclusion

Starting with the marginal risk-neutral distributions of the index and its expected volatility, we describe a way to identify the joint distribution. We achieve identification of the parameter of a Frank copula (which determines the correlation) by fitting two-month market index options based on its two-month risk-neutral distribution. This distribution we build up from the one-month return distribution and rescaled second-month distributions, which incorporate the information in the expected volatility distribution and the correlation parameter.

The implied index-to-volatility correlation premium turns out to be significantly negative and time-varying. It predicts future index returns above and beyond the variance risk premium. It is related to measures of risk and to the persistence of risk in a sub-additive way. That means that the implied index-to-volatility correlation premium (which we interpret to measure future investment opportunities and also investor attitude towards market down turns) moves more negative (worse outlook) as current risk is higher and risk persistence is higher. However, the

interaction term works the other direction, so high and persistent risk simultaneously leads to a better outlook that the sum of the two effects separately.

Our work provides new ideas for risk management and pricing of portfolios of index and/or volatility related securities. In the future, we intend to look at bivariate pricing kernels and asset pricing tests, which could use index-to-volatility correlation as factors.

A Technical Appendix

This section provides necessary details for the simulation-based calibration of the index-to-volatility implied correlation.

1. Preparation of data on S&P 500 and VIX options

   (a) Process IvyDB OptionMetrics data using raw options for S&P 500 and VIX. Raw data includes underlying spot price, bid/ask prices of options, mid-point implied volatility, date of observation, and days to maturity.

   (b) Recompute implied volatilities (IVs) for VIX options. The OptionMetrics database (version update on WRDS from March 2013) contains an error in computing implied volatility and Greeks of the VIX options—instead of using the futures level, OptionMetrics uses current VIX level as the at-the-money reference level without using the correct cost of carry/ convenience yield. Instead of using cost of carry/ convenience yield in the Black-Scholes formula one can use the futures level directly and infer the implied volatility by inverting Black (1976) formula.

   We merge options data with end-of-day data on futures (from CBOE) and compute the implied volatilities by inverting the Black formula:

   $$Call(T, F, K) = e^{-rT}[F \times N(d_1) - K \times N(d_2)],$$

   $$Put(T, F, K) = e^{-rT}[K \times N(-d_2) - F \times N(-d_1)],$$

   $$d_1 = \frac{\log(F/K) + (\sigma^2/2)T}{\sigma\sqrt{T}}, \ d_2 = d_1 - \sigma\sqrt{T},$$

   where T is the time to maturity, F is the futures level, K is the strike, r is the prevailing riskfree rate, $N(\cdot)$ is the cumulative standard normal distribution function, and σ is the implied volatility we are looking for.

   (c) For each month we take the regular *expiration* day on the third Friday and find the *observation* date in the previous month closest to 25 days before the selected

expiration day. We select all options observed on the *observation* date and expiring

on the regular third Friday expiration date in the next month for short-term options

and in the second month for longer term options. We use 25 days to maturity as the

target duration for short maturity options, because (i) it allows to select the option

sample with the most stable days to maturity each month, (ii) we obtain second

maturity (longer-term options) closest to 60 days, and (iii) for 25 days we compute

the estimated index-to-volatility implied correlation closest to the end of the month.

The typical date on which we perform the calibration is between the 20th and the

28th of each month.

(d) On each monthly calibration date, select all observed S&P 500 and VIX options the

next and the second month maturities. For each data point record the S&P 500

continuous dividend yield δ_t and interpolate the zero CD rate (from OptionMetrics)

to the exact days to maturity of the selected S&P 500 options to find the continuously

compounded rate $r_{f,t}$.

(e) Select the short-term options for calibration on a given date. For S&P 500 use out-

the-money (OTM) calls and puts with moneyness (defined as K/S with respect to

the current S&P 500 level S) between 0.7 and 1.3, i.e., filter out puts with $K/S > 1$

and calls with $K/S < 1$, and for both types of options filter out $K/S < 0.7$ and

$K/S > 1.3$.[8] For VIX use all OTM calls and puts, where moneyness is determined

with respect to the futures level (i.e., moneyness is K/F). For both underlying

securities we normalize option strike by the underlying price (i.e., spot price for

S&P 500, and futures level for VIX), so that each option strike now represents the

gross return (or relative change) for the given underlying on the maturity date. We

estimate the marginal RND from S&P 500 and VIX options for the dates, where we

have at least 8 and 5 available options, respectively.

---

[8]For robustness we also apply different filters on the moneyness of options.

2. Estimate marginal risk-neutral distributions (RND) for S&P 500 from observed short-term options using the "fast and stable method" of Jackwerth (2000): the method finds a smooth risk-neutral distribution that, at the same time, explains the option prices.[9]

   (a) Fit the implied volatility curve as a function of the moneyness with values σ_j at discrete moneyness grid points $m_j \in [0.5, 1.5]$, $j = 1 \ldots N$. To fit the implied volatility function we select the step size Δ and smoothness parameter λ for implied volatility interpolation to achieve RND, which is smooth. Because the "fast and stable method" leads to a system of equations, where coefficients are functions of the step size on a moneyness grid, and because S&P 500 options are too close to each other in moneyness, there is not enough machine accuracy to solve the system. To circumvent the problem, we multiply all moneyness points by 1000 and divide by the same multiplier when we need raw returns. Output from the procedure: a *grid of raw returns* (i.e., moneyness points) m_j between 0.5 and 1.5, with respective implied volatilities σ_j for all points on the grid. The risk-neutral probability distribution $q(m_j)$ can be written directly as

$$q(m_j) = e^{rT} \left[\frac{e^{-rT} n(d2_j)}{m_j \sigma_j \sqrt{T}} (1 + 2m_j d1_j \sigma_j' \sqrt{T}) + e^{-\delta T} n(d1_j) \sqrt{T} \left(\sigma_j'' + \frac{d1_j d2_j}{\sigma_j} (\sigma_j')^2 \right) \right],$$
(A1)

$$d1_j = \frac{-\log(m_j) + (r - \delta)T + (\sigma_j^2/2)T}{\sigma_j \sqrt{T}}, \quad d2_j = d1_j - \sigma_j^2/2T,$$

where T is the time to maturity, m_j is the moneyness level at grid point j, r is the prevailing riskfree rate, δ is the dividend yield, σ_j is the implied volatility at the grid point j, σ_j' and σ_j'' are the first and second derivatives of the volatility function with respect to moneyness at grid point j, and $n(\cdot)$ is the standard normal density function.

   (b) Due to requirements on the smoothness of the resulting risk-neutral distribution, the interpolated and extrapolated implied volatility curve does not exactly fit the

---

[9]See Appendix A in Jackwerth (2000) for technical details.

initially observed option prices and their implied volatilities. To improve the fit of the risk-neutral distribution to the observed option prices, while maintaining its smoothness controlled directly by the parameters of the fast and stable method, we introduce the second step in RND estimation. After deriving the "fast and stable" RND, we adjust it in two ways: (*) first, we set the mean of the distribution equal to the risk-free rate (demean first the resulting distribution from the previous step and add a new mean $= e^{r_{f,t}-\delta_t}$ (or $\frac{1+r_{f,t}}{1+\delta_t}$ as an approximation), adjusted for the time to maturity T:

$$m_j = m_j - \sum_{z=1}^{N} q_{m_z} \times m_z + e^{(r-\delta)T}, \quad \forall j = 1, \ldots, N, \tag{A2}$$

and (**) second, we set the volatility of the distribution to the value that provides the best fit of the option prices computed from the derived RND to the observed options prices. To price an option we compute the payoff of a given option for each grid point, compute the expectation of the payoff using the adjusted RND, and discount

at the riskfree rate.

$$\hat{\sigma}_0 = \sqrt{\sum_{z=1}^{N} q_{m_z} \times (m_z - 1)^2 - \left(\sum_{z=1}^{N} q_{m_z} \times (m_z - 1)\right)^2},$$

$$mean_0 = \sum_{z=1}^{N} q(m_z) \times m_z,$$

$$m_j = [m_j - mean_0] \times \frac{\hat{\sigma}_{0,opt}}{\hat{\sigma}_0} + mean_0, \quad \forall j = 1, \ldots, N,$$

$$Prc_c^{est} = e^{-rT} \sum_{z=1}^{N} q(m_z) \times \max(0, (m_z - 1) - \frac{K_c}{S}), \quad \forall c \in \{\text{Observed Calls}\},$$

$$Prc_p^{est} = e^{-rT} \sum_{z=1}^{N} q(m_z) \times \max(0, \frac{K_p}{S} - (m_z - 1)), \quad \forall p \in \{\text{Observed Puts}\},$$

$$RMSE(Prc) = \sqrt{\frac{1}{\#j \in \{c,p\}} \sum_{j \in \{c,p\}} \left(Prc_j^{est} - Prc_j^{obs}\right)^2},$$

$$\hat{\sigma}_{0,opt} = \arg\min RMSE(Prc). \tag{A3}$$

We guarantee that the sum of all risk-neutral probabilities q_j, $\forall j = 1, \ldots, N$ is equal to one, by normalizing each probability q_j by the sum of probabilities.

3. Estimate marginal RND for future expected volatility from observed short-term VIX options using the "fast and stable method."

   (a) Fit the implied volatility curve as a function of the moneyness with values σ_j at discrete moneyness grid points $m_j \in [0.2, 2.2]$, $j = 1 \ldots N$. Control the step size Δ and smoothness parameter λ for implied volatility interpolation to achieve RND, which is smooth. As in case with S&P 500, we multiply all moneyness points by 1000 and divide by the same multiplier when we need raw returns. Output from the procedure: a *grid of raw VIX returns respective to the futures level* (i.e., moneyness points) m_j between 0.2 and 2.2, with respective implied volatilities σ_j for all points

24

on the grid. The risk-neutral probability distribution $q(m_j)$ can be written directly as a function of implied volatilities, option moneyness and maturity, as shown for the case of S&P 500 in the equation (A1); note that because VIX is not traded, we set both interest rate r and dividend yield δ to zero, using moneyness m_j defined with respect to the VIX futures level.

(b) Due to requirements on the smoothness of the resulting risk-neutral distribution, the interpolated and extrapolated implied volatility curve does not exactly fit the initially observed option prices and their implied volatilities. To improve the fit of the risk-neutral distribution to the observed option prices, while maintaining its smoothness controlled directly by the parameters of the fast and stable method, we introduce the second step in RND estimation, where we adjust the volatility of the distribution so that it obtains the best fit of the option prices computed from the derived RND to the observed options prices. As opposed to S&P 500 RND estimation, where we apply two adjustments—mean as in the equation (A2), and volatility as in the system (A3), the drift of the risk-neutral distribution of VIX is not fixed at the level of riskfree rate, and hence we apply only the volatility adjustment; the mean of the VIX changes RND is just equal to one, because the changes are defined relative to VIX futures level. To price the options we compute the payoff of a given option for each grid point, and then compute the expectation of the payoff using the adjusted RND. The procedure is equivalent to the one applied to S&P 500 distribution in system (A3).

4. Generate a two-period (approximately 60 days, which correspond to the maturity of a second-maturity options on each calibration date) S&P 500 risk-neutral distribution.

   (a) On each calibration day we are given: one-period risk-neutral density q^{SP} for S&P 500 with the support $m_j^{SP} \in [0.5, 1.5]$, $j = 1 \ldots N^{SP}$, one-period risk-neutral density q^{VIX} for VIX changes from its reference value with the support $m_j^{VIX} \in$

$[0.2, 2.2]$, $j = 1 \ldots N^{VIX}$. We use each density function and the grid points from the respective support to estimate the interpolated inverse cumulative distribution function $Qinv^{SP}$ and $Qinv^{VIX}$, s.t.

$$m^{SP} = Qinv^{SP} \left(\int_{-\infty}^{m^{SP}} q^{SP}(m)dm \right),$$

$$m^{VIX} = Qinv^{VIX} \left(\int_{-\infty}^{m^{VIX}} q^{VIX}(m)dm \right).$$

We generate $10,000$ draws from a joint distribution of S&P 500 and VIX, where the joint density is defined by the Frank copula $C(\cdots)$ with a dependency parameter θ, and marginal distributions of S&P 500 and VIX. Each draw is given by a pair of uniformly distributed numbers u^{SP}, $u^{VIX} \sim U[0,1]$, where u^{SP} is iid, and u^{VIX} is drawn conditionally on u^{SP} and copula parameter θ. The corresponding values of S&P 500 and VIX realizations are given by $m^{SP} = Qinv^{SP}(u_1)$ and $m^{VIX} = Qinv^{VIX}(u_2)$, respectively. The risk-neutral joint cumulative distribution function for a given pair of realizations is given by the copula function:

$$Q(m^{SP}, m^{VIX}) = Q\left(Qinv^{SP}(u^{SP}), Qinv^{VIX}(u^{VIX})\right) = C(u^{SP}, u^{VIX}; \theta).$$

(b) For each pairwise realization $i = 1, \ldots, 10,000$ of S&P 500 return m_i^{SP} and expected volatility change m_i^{VIX} construct the conditional second-period distribution of S&P 500 return, such that its volatility is equal to the drawn realization of the future expected volatility. The distribution will have the same probabilities q_j^{SP} of each realization of the gross return for a grid point j as in the first period, but the grid point value will change from m_j^{SP} to $m_{j,1}^{SP}$ to reflect the volatility realization drawn for this run. To adjust for new volatility, we proceed in three steps: first, we demean the first-period S&P 500 return distribution:

$$m_j^{SP,demeaned} = m_j^{SP} - \sum_{z=1}^{N^{SP}} q^{SP}(m_z) \times m_z^{SP}, \ \forall j = 1, \ldots, N^{SP},$$

second, we normalize all support points m_j^{SP} by the volatility $\hat{\sigma}_{0,opt}$, which we estimated by solving the problem described by the system (A3), and multiply all support

points by the volatility $\sigma_{1,i} = m_i^{VIX} E[\sigma_1]$, derived from the realization of volatility change m_i^{VIX} for a given draw i and future expected volatility level $E[\sigma_1]$, estimated jointly with the copula parameter θ as described below, and third, we add the mean of the initial S&P 500 distribution back. The changes in the mean and volatility of the distribution are certainly adjusted to reflect the differences in the length of the first T_0 and the second T_1 period (expressed as fractions of the year).

$$m_{j,1}^{SP} = m_j^{SP,demeaned} \times \frac{m_i^{VIX} E[\sigma_1]}{\hat{\sigma}_{0,opt}} + \sum_{z=1}^{N^{SP}} q^{SP}(m_z) \times m_z^{SP}, \forall j = 1, \ldots, N^{SP} \quad (A4)$$

(c) Construct $10,000$ two-period distributions of S&P 500 gross return for respective draws of S&P 500 return and VIX. The support of each two-period distribution of S&P 500 gross return conditional on the draw $\{m_i^{SP}, m_i^{VIX}\}$ is given by the product of the first period draw of the gross return realization m_i^{SP} and each adjusted realization of second-period S&P 500 gross return $m_{j,1}^{SP}$ from (A4):

$$m_{j,i,2}^{SP} = m_i^{SP} \times m_{j,1}^{SP}, \ \forall j = 1, \ldots, N^{SP}.$$

Each of such conditional distributions has the same probability density function $q_i^{SP}, \ \forall j = 1, \ldots, N^{SP}$, but a different support depending on the realization of the random copula-based draw. We can stop here and use $10,000$ resulting conditional distributions in a Monte-Carlo procedure to price the 2-period (60-day) options on S&P 500. Because we want to analyze the properties of the two-period and of the second-period distributions, we go one step further and construct an unconditional two-period distribution of S&P 500 gross return with the support consisting of 500 realizations of two-period return.

We proceed as follows: (i) collect all 10,000 realizations m_i^{SP} of the first-period S&P 500 gross return from the copula simulation procedure and allocate them to the 500 equally spaced bins within the limits that cover all realizations; find the probabilities of generating each pairwise draw u^{SP}, u^{VIX} from the copula simulation procedure (i.e., from the *copulapdf* in Matlab); and normalize the probabilities

27

of the return draws falling into each bin so that within each bin the probabilities of all observations sum up to one. Multiply the probability of each draw in a bin by the probability of observing a given S&P 500 draw obtained from the marginal S&P 500 distribution density q^{SP} to get the probability of observing each S&P 500 return in the first period and a drawn value of the future expected volatility.

(ii) collect all the realizations of the conditional two-period distributions $m_{j,i,2}^{SP}, \forall j = 1, \ldots, N^{SP}$, $i = 1, \ldots, 10,000$ and allocate them to 500 equally spaced bins, truncating any realizations below 0 and above 2; from the allocated realizations compute the average realization (gross return) $m_{2,j}^{SP}$ in each bin $j = 1, \ldots, 500$, and allocate to this moneyness a probability density $q_{2,j}^{SP}$ equal to the sum of the product of the probabilities of the first-period S&P 500 return realizations and probabilities of the second-period realizations from each conditional distribution, such that the pair of first- and second-period returns lead to the given two-period S&P 500 realizations in a given bin.

(d) Use the two-period S&P 500 distribution to price two-period (approximately 60-day) S&P 500 options. Before using the RND to price options, adjust the two-period S&P 500 return distribution the same way as we did in equation (A2) for the first-period RND, i.e., set its drift equal to the $e^{(r_f - \delta)T_2}$, where T_2 is the length of the two periods. Then compute the payoff for an option with a given strike price and maturity for each of the 500 grid points, compute the expectation of the payoff using the adjusted RND, and discount at the riskfree rate:

$$Prc_c^{est} = e^{-rT_2} \sum_{z=1}^{500} q_2^{SP}(m_{2,z}^{SP}) \times ((m_{2,z}^{SP} - 1) - \frac{K_c}{S})^+, \forall c \in \{\text{Obserbed Calls}\}, \text{(A5)}$$

$$Prc_p^{est} = e^{-rT_2} \sum_{z=1}^{500} q_2^{SP}(m_{2,z}^{SP}) \times (\frac{K_p}{S} - (m_{2,z}^{SP} - 1))^+, \forall p \in \{\text{Observed Puts}\}. \text{(A6)}$$

(e) Invert the Black and Scholes (1973) option pricing formula to get the implied volatilities from the observed two-period option prices IV^{obs}, and from the option prices

for the same strikes estimated in equations (A5) and (A6), i.e., IV^{est}. Calibrate the copula parameter θ and future expected volatility level $E[\sigma_1]$ to minimize the pricing error in terms of implied volatilities in two-period (60-day) S&P 500 options (we minimize the RMSE of IV^{est} relative to the observed IV^{est} for all available out-the-money options within moneyness range of [0.8,1.2]):

$$\{\theta, E[\sigma_1]\} = \arg\min \sqrt{\frac{1}{\#j \in \{c,p\}} \sum_{j \in \{c,p\}} \left(IV_j^{est} - IV_j^{obs}\right)^2}.$$

5. Estimate the index-to-volatility correlation directly from drawn 10,000 pairwise realizations of S&P 500 gross return and VIX changes.

References

Allen, David E., Michael McAleer, Robert Powell, and Abhay K. Singh, 2012, A Non-Parametric and Entropy based Analysis of the Relationship between the VIV and S&P 500, Working Paper, Edith Cowan University.

Andersen, Torben G., Tim Bollerslev, Francis X. Diebold, and Heiko Ebens, 2001, The Distribution of Stock Return Volatility, Journal of Financial Economics 61, No. 1, 43-76.

Avellaneda, Marco, 2013, New Techniques for Pricing VIX Futures and VXX Options, Working Paper, New York University.

Bakshi, Gurdip, Dilip B. Madan, and George Panayotov, 2010, What do VIX Puts Reveal about the Volatility Market? Working Paper, University of Maryland.

Bakshi, Gurdip., Nikunj Kapadia, and Dilip B. Madan, 2003, Stock Return Characteristics, Skew Laws, and the Differential Pricing of Individual Equity Options, The Review of Financial Studies 16, No. 1, 101-143.

Baltussen, Guido, Sjoerd Van Bekkum, and Bart Van Der Grient, 2012, Unknown Unknowns: Vol-of-Vol and the Cross Section of Stock Returns, Working Paper, Erasmus University Rotterdam.

Bardgett, Chris, Elise Gourier, and Markus Leipold, 2013, Inferring Volatility Dynamics and Risk Premia from the S&P 500 and VIX Markets, Working Paper, University of Zurich.

Barndorff-Nielsen, Ole E., and Neil Shephard, 2002, Econometric Analysis of Realized Volatility and Its Use in Estimating Stochastic Volatility Models, Journal of the Royal Statistical, Society Series B 64, No. 2, 253-280.

Bauwens, Luc, Sebastian Laurent, and Jeroen V. K. Rombouts, 2006, Multivariate GARCH Models: A Survey, Journal of Applied Econometrics 21, No. 1, 79-109.

Bekaert, Geert, and Guojun Wu, 2000, Asymmetric Volatility and Risk in Equity Markets, Review of Financial Studies 13, No. 1, 1-42.

Bekaert, Geert, and Marie Hoerova, 2013, The VIX, the Variance Premium and Stock Market Volatility, Working Paper, Columbia University.

Black, Fischer, and Myron Scholes, 1973, The Pricing of Options and Corporate Liabilities, Journal of Political Economy 81, No. 3, 637-654.

Black, Fisher, 1976, The Pricing of Commodity Contracts, Journal of Financial Economics 3, No. 1-2, 167-179.

Bollerslev, Tim, George Tauchen, and Hao Zhou, 2009, Expected Stock Returns and Variance Risk Premia, Review of Financial Studies 22, No. 11, 4463-4492

Bollerslev, Tim, and Viktor Todorov, 2011, Tails, Fears, and Risk Premia, Journal of Finance 66, No. 6, 2165-2211.

Branger, Nicole, and Clemens Voelkert, 2013, The Fine Structure of Variance: Consistent Pricing of VIX Derivatives, Working Paper, Westfaelische Wilhelms Universitaet Muenster.

Carr, Peter, and Dilip B. Madan, 2013, Joint Modeling of VIX and SPX Options at a Single and Common Maturity with Risk Management Applications, Working Paper, New York University.

Carr, Peter, and Liuren Wu, 2009, Variance Risk Premiums, Review of Financial Studies 22, No. 3, 1311-1341.

Chabi-Yo, Fousseni, 2012, Pricing Kernels with Stochastic Skewness and Volatility Risk. Management Science 58, No. 3, 624–640.

Chabi-Yo, Fousseni, and Zhaogang Song, 2013, Recovering the Probability Weights of Tail Events with Volatility Risk from Option Prices, Working Paper, Ohio State University.

Chang, Bo Young, Peter Christoffersen, and Kris Jacobs, 2012, Market Skewness Risk and the Cross Section of Stock Returns, Journal of Financial Economics 107, No. 1, 46-68.

Christoffersen, Peter, Kris Jacobs, and Bo Young Chang, 2012, Forecasting with Option-Implied Information, Handbook of Economic Forecasting 2, Graham Elliott and Allan Timmermann.

Christoffersen, Peter, Steven Heston, and Kris Jacobs, 2010, Option Anomalies and the Pricing Kernel, Unpublished Working Paper, University of Maryland.

Cremers, Martijn, Michael Halling, and David Weinbaum, 2012, Aggregate Jump and Volatility Risk in the Cross-Section of Stock Returns, Working Paper, University of Notre Dame.

Driessen, Joost, Pascal J. Maenhout and Grigory Vilkov, 2013, Option-Implied Correlations and the Price of Correlation Risk, Working Paper, Advanced Risk & Portfolio Management Paper.

Duan, Jin-Chuan, and Chung-Ying Yeh, 2012, Price and Volatility Dynamics Implied by the VIX Term Structure, Working Paper, National University of Singapore.

Epps, Thomas W., 1979, Comovements in Stock Prices in the Very Short Run, Journal of the American Statistical Association 74, No. 366, 291-298.

Frank, Maurice J., 1979, On the Simultaneous Associativity of F(x,y) and x+y - F(x,y), Aequationes Mathematicae 19, No. 1, 194–226.

Fuertes, Carlos, and Andrew Papanicolaou, 2012, Implied Risk-Neutral Filtering Densities on Volatility's Hidden State, Working Paper, Princeton University.

Gabaix, Xavier, 2012, Variable Rare Disasters: An Exactly Solved Framework for Ten Puzzles in Macro-Finance, The Quarterly Journal of Economics 127, No. 12, 645-700.

Heston, Steven L., 1993, A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options, Review of Financial Studies 6, No. 2, 327-343.

Ishida, Isao, Michael McAleer, and Kosuke Oya, 2011, Estimating the Leverage Parameter of Continuous-Time Stochastic Volatility Models Using High Frequency S&P 500 and VIX, Working Paper, Osaka University.

Jackwerth, Jens Carsten, 2004, Option-Implied Risk-Neutral Distributions and Risk Aversion, Research Foundation of AIMIR, CFA Institute-Publications.

Javier, Mencia, and Enrique Sentana, 2013, Valuation of VIX Derivatives, Journal of Financial Economics 108, No. 2, 367-391.

Khagleeva, Inna, 2012, Understanding Jumps in the High-Frequency VIX, Working Paper, University of Illinois at Chicago.

Martin, Ian, 2013, Simple Variance Swaps, Working Paper, Stanford GSB and NBER.

Patton, Andrew, 2009, Copula-Based Models for Financial Time Series, Handbook of Financial Time Series, Springer Verlag, 767-785.

Shen, Xiaoyu, 2013, Two Volatility Term Structures, Working Paper, Tinbergen Institute and University Amsterdam.

Song, Zhaogang, and Dacheng Xiu, 2013, A Tale of Two Option Markets: Pricing Kernels and Volatility Risk, Working Paper, University of Chicago.

Trevedi, Pravin K., and David M. Zimmer, 2005, Copula Modeling: An Introduction for Practitioners, Econometrics 1, No. 1, 1-111.

Whaley, Robert E., 2013, Trading Volatility: At What Cost?, Working Paper, Vanderbilt University.

Zhang, Lan, Per A Mykland, and Yacine Ait-Sahalia, 2005, A Tale of Two Time Scales: Determining Integrated Volatility with Noisy High-Frequency Data, Journal of the American Statistical Association 100, No. 472, 1394-1411.

Table 1: Explanations of tail fatness and pricing kernel shifts.

In Panel A, we define our dependent variable of tail fatness to be the ratio of left tail probability beyond the 10th percentile return (-7.77% monthly in our sample) for the second period over the first period. In Panel B, we define the slope of the ratio of first period over second period pricing kernels (shift in pricing kernels). Assuming that the actual distribution does not change from one month to the next, we compute the ratio of the pricing kernels as the ratio of the second months to the first months risk-neutral distribution. We compute the slope between the 5th and the 30th percentile to be our $shift_t$ variable. We regress tail fatness on a constant, the index-to-volatility correlation risk premium, the variance risk premium, and the implied skewness of the first period risk-neutral distribution.

Panel A: Tail Fatness

	(1)	(2)	(3)
Intercept	1.1179	1.1316	0.8363
p-val	0.0000	0.0000	0.0000
Index-to-vol Corr Prem	-0.4001	-	-0.4349
p-val	0.0000	-	0.0000
Variance Risk Premium	-	1.2271	1.1317
p-val	-	0.2303	0.1069
Implied Skewness, 30d	-	-	-0.1282
p-val	-	-	0.0002
Rbar	0.1558	0.0241	0.3635

Panel B: Shift in Pricing Kernels

	(1)	(2)	(3)
Intercept	-1.0231	-1.1097	1.0059
p-val	0.0000	0.0000	0.0677
Index-to-vol Corr Prem	3.0225		3.2787
p-val	0.0000		0.0000
Variance Risk Premium		-10.4934	-9.8529
p-val		0.1785	0.0715
Implied Skewness, 30d			0.9117
p-val			0.0004
Rbar	0.1593	0.0354	0.3619

Table 2: Return predictability.

We define our dependent variable as the 25 day return on the S&P 500, i.e., the return from the day of calibration until the maturity of the nearest maturity options used in calibration. We regress future returns on a constant, the index-to-volatility correlation risk premium, the variance risk premium, and the implied skewness of the first period risk-neutral distribution.

	(1)	(2)	(3)
Intercept	-0.0016	-0.0047	-0.0212
p-val	0.7937	0.4521	0.3581
Index-to-vol Corr Prem	-0.0636	-	-0.0676
p-val	0.0176	-	0.0181
Variance Risk Premium	-	0.5692	0.5753
p-val	-	0.0098	0.0027
Implied Skewness, 30d	-	-	-0.0054
p-val	-	-	0.5560
Rbar	0.0384	0.0822	0.1194

**Table 3:** Explaining implied risk-neutral index-to-volatility correlation.

Our dependent variable is the index-to-volatility correlation premium. Independent variables are autocorrelation in VIX, historical volatility, VIX, the volatility of historical variance, and option implied correlation among the equities of the S&P 500. We also use interacted variables where we multiply VIX autocorrelation with eac of the remaining risk measures.

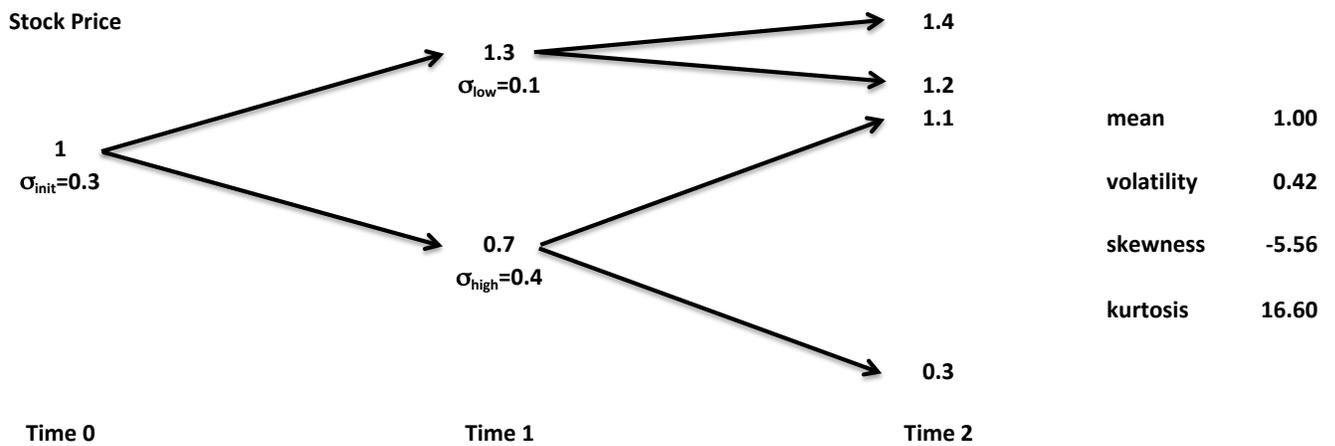|	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Intercept	0.1124	-0.0915	0.3920	-0.1143	0.4589	-0.0822	0.1829	-0.1325	0.4719
p-val	0.2594	0.0119	0.0116	0.0046	0.0099	0.0003	0.0907	0.0202	0.0364
VIX Autocorrelation	-0.1045	-	-0.2663	-	-0.3205	-	-0.1450	-	-0.3446
p-val	0.0427	-	0.0011	-	0.0009	-	0.0104	-	0.0051
Past SP500 Volatility	-	0.0727	-1.4741	-	-	-	-	-	-
p-val	-	0.5432	0.0064	-	-	-	-	-	-
VIX	-	-	-	0.1582	-1.3798	-	-	-	-
p-val	-	-	-	0.2343	0.0145	-	-	-	-
Vol-of-Var	-	-	-	-	-	0.2952	-4.3708	-	-
p-val	-	-	-	-	-	0.1520	0.0307	-	-
SP500 Implied Correlation	-	-	-	-	-	-	-	0.1227	-0.7687
p-val	-	-	-	-	-	-	-	0.2828	0.0896
Past SP500 Volatility * AC	-	-	0.8267	-	-	-	-	-	-
p-val	-	-	0.0038	-	-	-	-	-	-
VIX * AC	-	-	-	-	0.8538	-	-	-	-
p-val	-	-	-	-	0.0049	-	-	-	-
Vol-of-Var * AC	-	-	-	-	-	-	2.3546	-	-
p-val	-	-	-	-	-	-	0.0161	-	-
SP500 Implied Correlation * AC	-	-	-	-	-	-	-	-	0.5120
p-val	-	-	-	-	-	-	-	-	0.0290
Rbar	0.0434	-0.0102	0.0976	-0.0040	0.0914	-0.0094	0.0483	-0.0035	0.0670

Figure 1: The bivariate dependence of returns and volatility for different correlation levels.

We depict in Panel A a two-period binomial tree where volatility and returns have perfect positive correlation. The high (low) return is followed by a second period high (low) volatility return. In Panel B, the correlation is perfectly negative. Moments are for the two period distributions.

Panel A   Positive correlation index/volatility, rho = 1

Stock Price

mean	1.00
volatility	0.42
skewness	5.56
kurtosis	16.60
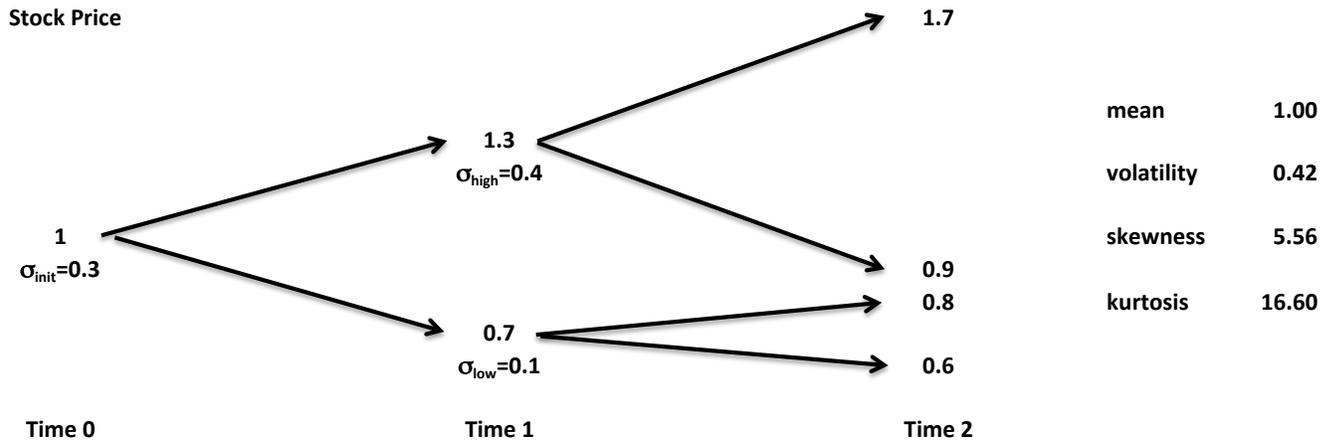
Panel B   Negative correlation index/volatility, rho = -1

Stock Price

mean	1.00
volatility	0.42
skewness	-5.56
kurtosis	16.60

Note: all move probabilities are 0.5

Figure 2: Risk-neutral distributions from one-month S&P500 and VIX options.

For one observation date in our sample (August 23, 2011), we show in the first row the observed vs. fitted implied volatilities of out-of-the-money options on S&P 500 and on VIX. In the second row we show the resulting risk-neutral distributions. Finally, third row depicts the out-of-the-money option prices (bid and ask) and the models based prices which lie in between.
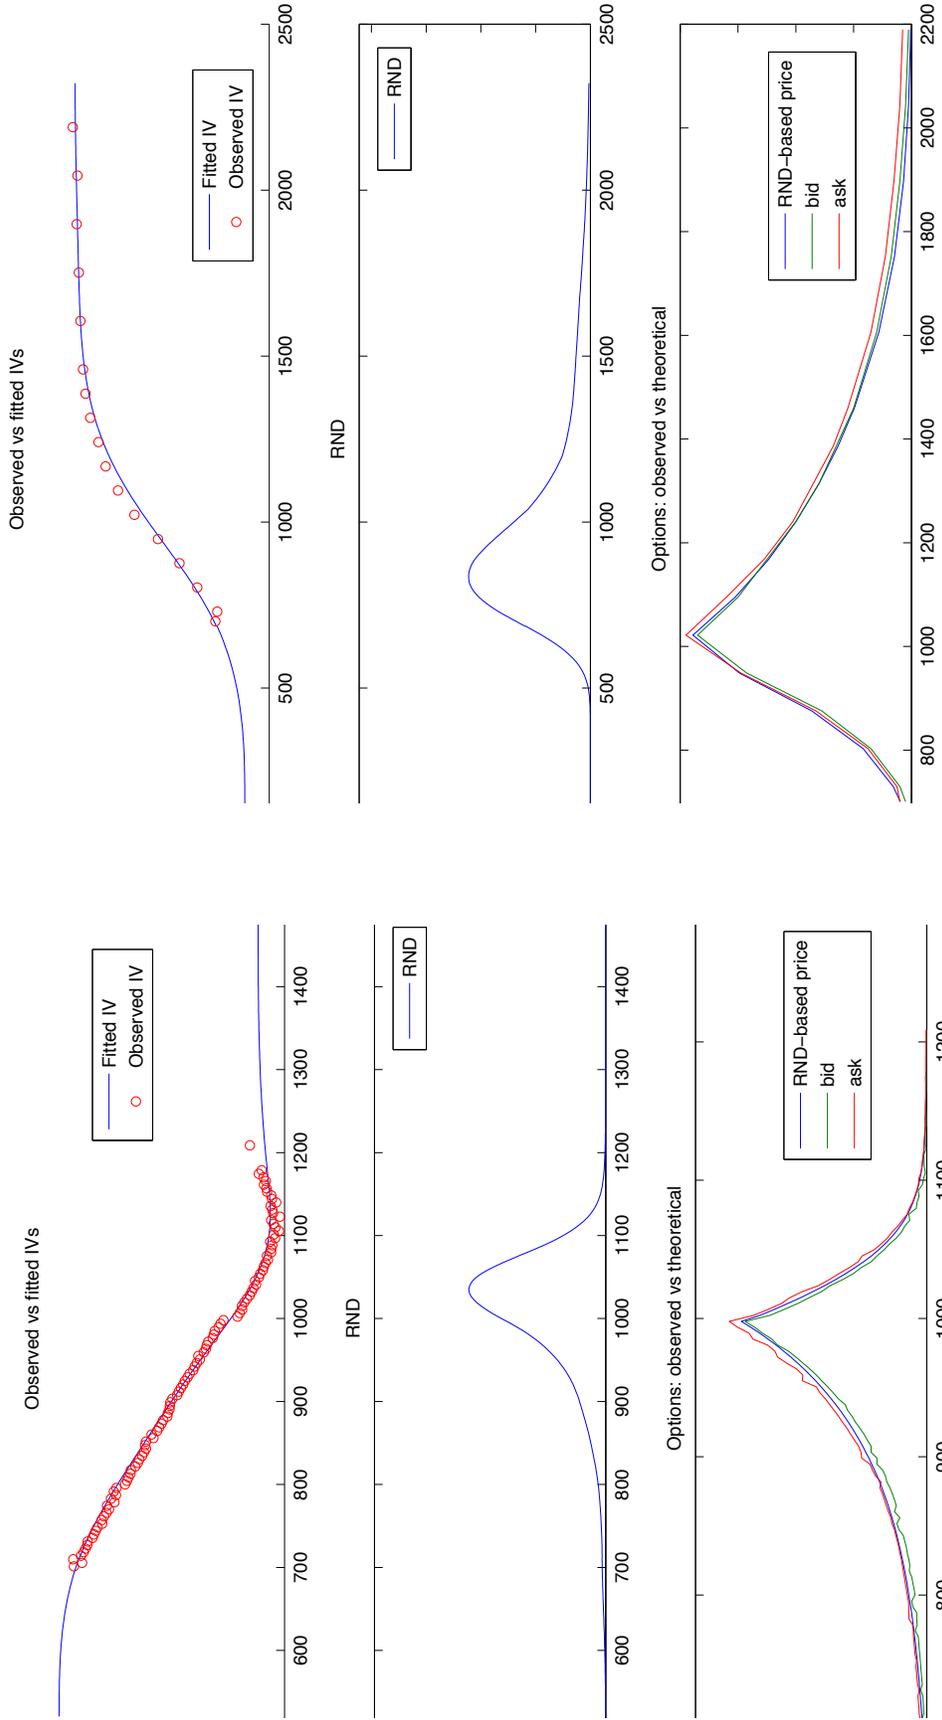
Figure 3: Bivariate distribution of one-month S&P500 returns and volatilities.

For one observation date in our sample (August 23, 2011), we show the marginal risk-neutral distributions of S&P 500 returns and volatilities. Using a Frank (1979) copula with correlation -0.87 between returns and volatilities, we depict the bivariate density as well. The parameter of the Frank copula and the mean of the second period volatility have been chosen so that the root mean error of the implied volatilities for the two month index options (model vs. observed) has been minimized.
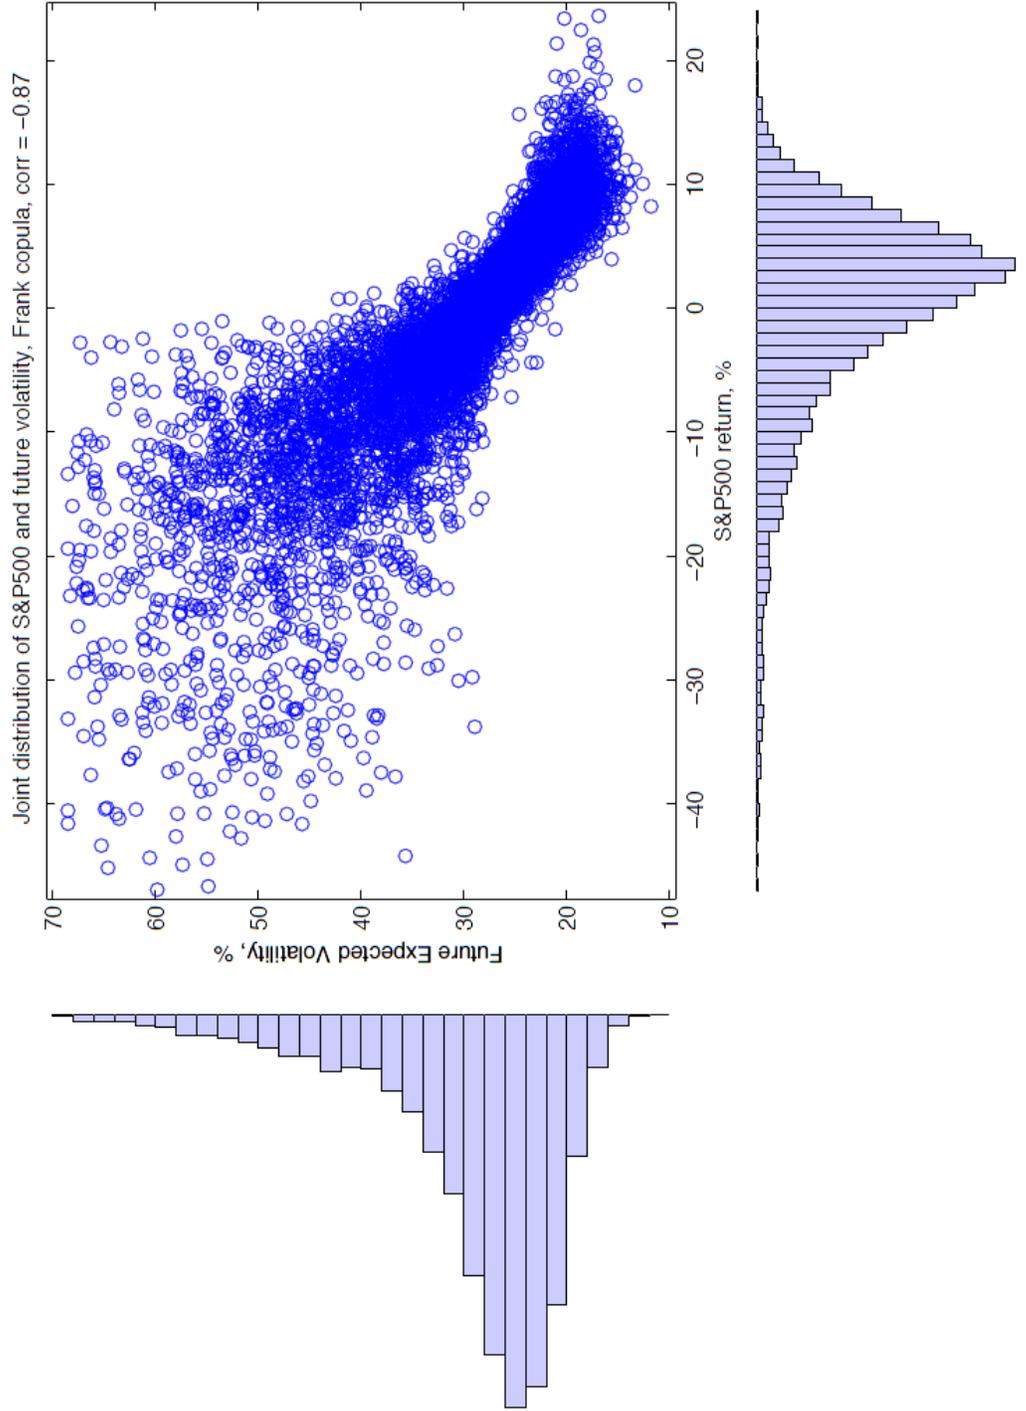
Figure 4: Prices of two-month S&P500 index options.

For one observation date in our sample (August 23, 2011), we show the (bid and ask) prices of two-month S&P 500 index options. We use a Frank (1979) copula with correlation -0.87 between returns and volatilities. The parameter of the Frank copula and the mean of the second period volatility have been chosen so that the root mean error of the implied volatilities for the two month index options (model vs. observed) has been minimized. We also show the model prices which lie in between the observed bid and ask prices.

dt: 89 date: 23–Aug–2011: 2–m opt, corr: –0.87279 copula: Frank