Asymmetric Volatility Risk: Evidence from Option Markets*

Jens Jackwerth **Grigory Vilkov**

This version: July 3, 2018

Abstract

Asymmetric volatility concerns the relation of returns to future expected volatility. Much is known from option prices about the marginal risk-neutral distributions of S&P 500 returns and of relative changes in future expected volatility (VIX). While the bivariate risk-neutral distribution cannot be inferred from the marginals, we propose a novel identification based on long-dated index options. We estimate the risk-neutral asymmetric volatility implied correlation and find it to be significantly lower than its realized counterpart. We interpret the economics of the asymmetric volatility correlation risk premium and use asymmetric volatility implied correlation to predict returns, volatility, and risk-neutral quantities.

Keywords: Asymmetric volatility, SPX options, VIX options, asymmetric volatility implied correlation, leverage effect, trading strategy

JEL: G11, G12, G13, G17

*Corresponding author: Jens Jackwerth is from the University of Konstanz, PO Box 134, 78457 Konstanz, Germany, Tel.: +49-(0)7531-88-2196, Fax: +49-(0)7531-88-3120, jens.jackwerth@uni-konstanz.de; Grigory Vilkov is from the Frankfurt School of Finance & Management, Germany, vilkov@vilkov.net.
We received helpful comments and suggestions from Raman Uppal, Bjorn Eraker, Axel Kind, Turan Bali, Guenter Franke, Nikunj Kapadia, Pascal Létourneau, Paulo Rodrigues, Loriana Pelizzon, and Paul Schneider. We thank participants of the Western Finance Association 2015 Meeting, IFSID Third Conference on Derivatives in Montreal, ESSFM 2013 Seminar at Gerzensee, and seminar participants at the Frankfurt School of Finance & Management, Goethe University of Frankfurt, University of Lugano, University of Manchester, University of Mannheim, the Swiss Finance Institute at EPFL, Tinbergen Institute/Duisenburg School of Finance, 2013 OptionMetrics Research Conference, University of Konstanz, EDHEC, University of South Australia, and University of Zurich.

Asymmetric Volatility Risk: Evidence from Option Markets

Abstract

Asymmetric volatility concerns the relation of returns to future expected volatility. Much is known from option prices about the marginal risk-neutral distributions of S&P 500 returns and of relative changes in future expected volatility (VIX). While the bivariate risk-neutral distribution cannot be inferred from the marginals, we propose a novel identification based on long-dated index options. We estimate the risk-neutral asymmetric volatility implied correlation and find it to be significantly lower than its realized counterpart. We interpret the economics of the asymmetric volatility correlation risk premium and use asymmetric volatility implied correlation to predict returns, volatility, and risk-neutral quantities.

1. Introduction

Asymmetric volatility is one of the fundamental drivers of asset prices in that returns are typically negatively correlated with future expected volatility. Using S&P 500 returns and relative changes in future expected volatility (as measured by VIX futures[1]), the average asymmetric volatility realized correlation (AVRC) of -0.77 indicates severe concerns of market participants about poor (good) returns that translate into elevated (subdued) levels of future expected volatility. The high empirical standard deviation of 0.10 and significant correlation of changes in AVRC with changes in VIX suggests that this risk perception might be time-varying.

Our goal is to understand which role asymmetric volatility plays in asset pricing. Thus, we are interested in the risk-neutral equivalent to AVRC, namely the asymmetric volatility implied correlation (AVIC) between returns and future expected volatility.[2] Much can be learned from short-dated option prices on the S&P 500 index and on VIX futures. The first set of options allows us to back out the risk-neutral distribution (RND) of index returns. The second set allows us to back out the RND of relative changes in future expected volatility. Yet, knowing the marginal RNDs does not tell us about the bivariate RND or about AVIC. Past research proceeded at this juncture by identifying implied correlation from basket or exchange options, written on all marginal quantities simultaneously.[3] Unfortunately, there are no basket options on returns and future expected volatility. Still there is a way to proceed.

We present a novel method for finding the bivariate RND in which we combine information from short-dated index and VIX futures options with information from long-dated index options. The basic idea is as follows. Assume that the index can only move up or down in the first period. Using a simple example in Figure 1, the index can move from

[1] VIX is an option-implied volatility index based on S&P 500 out-of-the-money options.

[2] We stay agnostic about the source of such correlation and it could come from either diffusive or jump components of the underlying stochastic processes.

[3] Driessen, Maenhout, and Vilkov (2005) use index options as basket options on stocks to find the implied correlation across stocks; Amato and Gyntelberg (2005) investigate CDO markets; Augustin et al. (2014) review research on CDS; and Mueller, Stathopoulos, and Vedolin (2017) look into FX markets.

1 up to 1.3 or down to 0.7 with equal probability, implying a first-period volatility of 0.3. Such an RND could be obtained from short-dated index options. From short-dated options on VIX futures, we learn that second-period volatility can increase to 0.4 or decrease to 0.1 with equal probability. Yet, we still do not know AVIC.

Assuming perfect *positive* AVIC in Panel A, volatility will increase to 0.4 after a first-period high return and volatility will be very low at 0.1 after a first-period low return. Therefore, the second-period returns should reflect this conditionality. We keep the first-period up and down probabilities in place during the second period at 0.5 and 0.5; that is, we assume that the shape of the conditional short-dated return RND remains the same in every period, and that its volatility can vary according to the RND of future expected volatility. Combining the first- and second-period returns, the two-period distribution will be more spread out for high returns (due to the high volatility after a first-period high return) compared to low returns. Assuming a zero interest rate and dividend yield for simplicity, we price a two-period call option with strike price 1.2 at 0.125.

Given perfect *negative* correlation in Panel B, the argument reverses exactly, and two-period returns will now be more spread out after a low first-period return and less spread out after a high first-period return. We price the same two-period call with strike price 1.2, but now its value is considerably lower at 0.05. Note that we can learn about AVIC from the price of that two-period call.

In Panel C, we assume that the call is priced at 0.0875. Allowing for intermediate correlations results in more complex mixtures of first- and second-period return distributions, which again yield different two-period return RNDs. We depict the situation with an AVIC of 0, which generates exactly the observed call price. Thus, we can use prices of two-period options on the S&P 500 to identify the AVIC and, more generally, the bivariate distribution RND between returns and future expected volatility.

In our empirical work, we reverse-engineer the AVIC by fitting options based on the two-period model RND to observed two-month options on the S&P 500 from July 2007 to August 2014. To find the two-period model RND, we infer the marginal RND of

Figure 1: Illustration of Asymmetric Volatility

We depict the bivariate dependence of returns (based on index levels S_t at $t = 0, 1, 2$) and volatility (σ) for different correlation levels. We depict in Panel A a two-period binomial tree in which volatility and returns have perfect positive correlation. The first-period high (low) return is followed by a second-period high (low) volatility return distribution. In Panel B, the correlation is perfectly negative. In Panel C, the correlation is zero. All move probabilities are 0.5, and the chance of a high and a low volatility regime is also 0.5. Call prices are for the two-period distributions with terminal nodal probabilities q, index levels of S_2, and payoffs C_2, based on a strike price of 1.2, the interest rate and dividend yield are zero.

Panel A: Perfect Positive Correlation between Returns and Future Expected Volatility

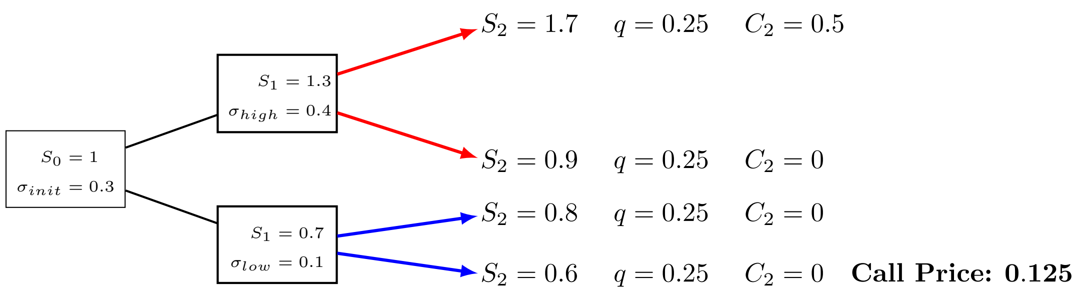

Panel B: Perfect Negative Correlation between Returns and Future Expected Volatility

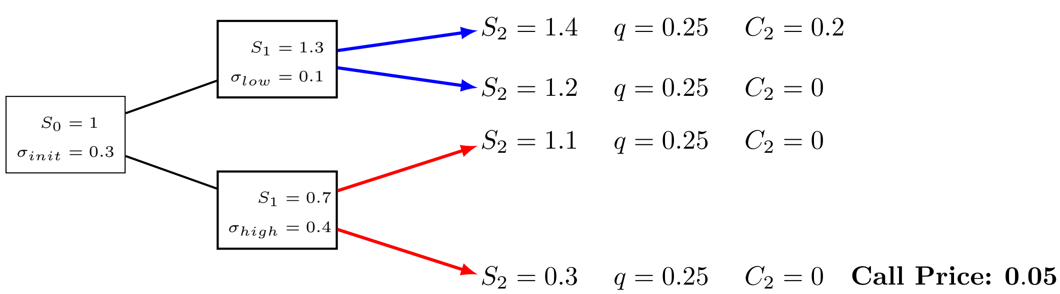

Panel C: Zero Correlation between Returns and Future Expected Volatility

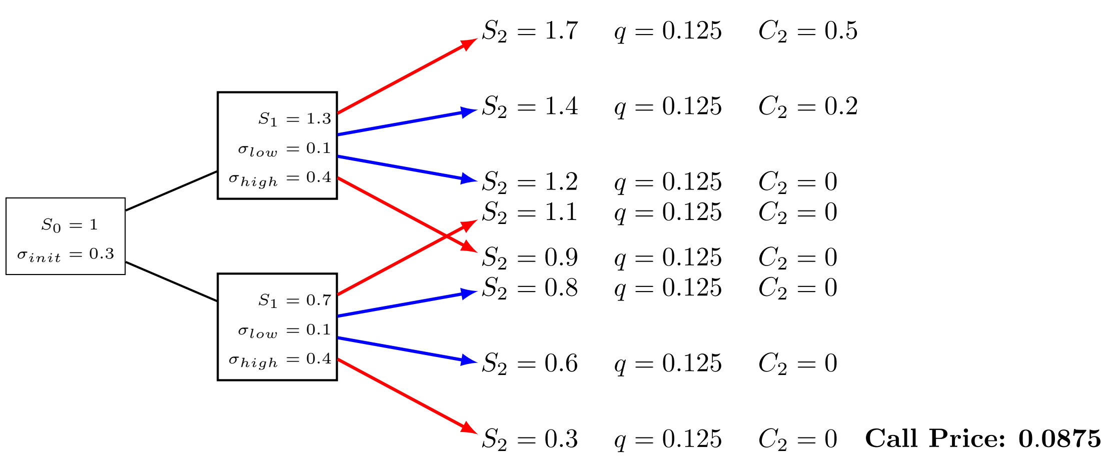

returns from one-month options on the S&P 500. We obtain the marginal RND of future expected volatility from one-month options on VIX futures.[4]

The dependence structure between the two marginal RNDs is modeled by a bivariate copula.[5] Each value of the dependency parameter of the copula (θ) maps into an AVIC. Given θ, we can build up the two-period risk-neutral return distribution, which we use to price two-period options. Finally, we vary θ to best fit the implied volatilities of the model options to the observed two-period option-implied volatilities. Details are collected in Section 4.

Empirically, we compute an AVIC every week during our sample from July 2007 to August 2014. The time-varying AVIC turns out to be -0.83 on average. We compute the AVRC based on a 360-day historical window, over which we correlate monthly returns with corresponding relative changes in VIX. The AVRC is at -0.77 significantly different from AVIC. Thus, there could be a positive asymmetric volatility correlation risk premium (AVCRP = AVRC-AVIC = 0.06) associated with asymmetric volatility risk.[6]

It is helpful to contrast the AVCRP with the variance risk premium (VRP). The VRP works as a symmetric fear gauge and measures the impact of variance risk in general. The AVCRP measures asymmetric volatility risk and is conditional on the first-period index return. We find that the AVCRP is largely driven by the AVIC. A more negative AVIC paired with a negative first-period return leads to a stronger increase in second-period volatility.

We concentrate on the AVIC and find that it predicts index returns, realized index volatility in some instances, and the *ex post* VRP. Also, the AVIC predicts risk-neutral quantities, such as fatter left tails of the RND, implied skewness in some instances, and

[4]Alternatively, one could have looked for options on future expected volatility under the physical measure, as opposed to options on VIX, which is the expected integrated volatility under the risk-neutral measure. Yet such alternative options do not exist.

[5]We use the Frank (1979) copula, which is particularly well-suited for the typically negative dependence we find in the data. It fits the data better than Gaussian and Student's t copulas, yet results are robust compared to using those alternative copulas; see Sections 4.3 and 6.

[6]Our AVCRP should not be confused with the risk premium for correlation between assets discussed in Buraschi, Trojani, and Vedolin (2014); Mueller, Stathopoulos, and Vedolin (2017); Buss, Schoenleber, and Vilkov (2018); and others.

implied volatility. Thus, a time-varying AVIC seems to be an important determinant of the future investment opportunities and should potentially be considered in modeling asset prices.

In terms of applications for our method of inferring the dependency between returns and future expected volatility, we know that their bivariate RND helps pricing contracts written on both quantities simultaneously. Also, term structures of the AVIC can be found from longer-dated options. Moreover, VIX-type contracts have recently been introduced in stock indices, individual stocks, commodity and country ETFs, interest rates, currencies, and volatilities. In 2011, the SEC was asked to permit options on 40 such contracts. Once they are available, our method could also be applied to all these new markets.

After putting our work into perspective with respect to the literature in Section 2, we introduce our data (Section 3) and methodology (Section 4). We present our results in Section 5 and show in Section 6 that they are robust to our choice of copula. We conclude in Section 7.

2. Literature

We use option and historical prices to relate returns to future expected volatility. This relates us directly to three large areas of research: first, the study of asymmetric volatility, which investigates the relation of returns and *future* expected volatility. Our paper uses this definition for all quantities concerned. Second is the literature concerning the relation of returns and *contemporaneous* volatility, which is often called the leverage effect.[7] Third is the study of recovering information implied in option prices.

[7] We thank a referee for pushing us to make this clarification, as the terminology in the literature at times deviates from our labeling.

2.1 Lead-lag "Asymmetric Volatility" Effect

Our work relates returns to future expected volatility. We label this lead-lag structure asymmetric volatility. Black (1976b) and Christie (1982) were the first to study stock returns and subsequent changes in stock return volatility.[8] Their empirical studies motivated discrete-time models with asymmetric volatility, such as the EGARCH model of Nelson (1991) and the asymmetric GARCH model of Glosten, Jagannathan, and Runkle (1993). These GARCH models assume constant parameters. We contribute to this literature by showing that asymmetric volatility is time-varying and different under the physical and risk-neutral probability measures.

Later papers try to explain the asymmetric volatility effect, in which Bekaert and Wu (2000) empirically use realized volatility, while Dennis, Mayhew, and Stivers (2006) use implied volatility. Figlewski and Wang (2000) question whether asymmetric volatility is really driven by the firm's leverage, and Hasanhodzic and Lo (2011) look at all-equity-financed firms to conclude that asymmetric volatility is not driven by financial leverage. We stay agnostic about the causes of asymmetric volatility but show that its time-varying nature constitutes a source of priced risk for investors.

2.2 Contemporaneous "Leverage" Effect

While discrete models often use a lead-lag structure of asymmetric volatility, many continuous-time models work with constant instantaneous correlation between returns and volatility, which we label the leverage effect (e.g., Bardgett, Gourier, and Leippold, 2017; Branger, Kraftschik, and Völkert, 2014; Christoffersen, Heston, and Jacobs, 2009; and Song and Xiu, 2016). Then, by Girsanov's theorem, the AVIC needs to equal the AVRC, and any AVCRP is necessarily zero, in contrast to our empirical findings. This rigidity in modeling is consistent with recent research documenting that the options on VIX futures cannot be priced properly using continuous-time models solely fitted to

[8] Confusingly, even though they use a lead-lag structure, they used the term "leverage effect," which we reserve for the contemporaneous relation of returns with volatility.

S&P 500 index options (e.g., Fuertes and Papanicolaou, 2014; Bates, 2012; Duan and Yeh, 2012; and Carr and Madan, 2013).

Andersen, Bondarenko, and Gonzalez-Perez (2015) review the problems of estimating the correlation between index returns and contemporaneous volatility (realized as well as a version based on a corridor-version of VIX) using high-frequency data. Kalnina and Xiu (2015) provide a theoretical base for such estimators. Aït-Sahalia, Fan, and Li (2013) concentrate on econometric issues and biases in estimating the leverage effect; Wang and Mykland (2014) estimate the leverage effect using high-frequency data; and Aït-Sahalia et al. (2017) split the leverage effect into continuous and discontinuous parts.

2.3 Recovering Information Implied in Market Prices

We recover the AVIC by semi-parametrically modeling the bivariate distribution between returns and future expected volatility. We thereby touch on recovering separate marginal RNDs of S&P 500 returns and relative changes in VIX futures. While we go on to analyze the joint distribution, there is a large body of literature on the extraction of marginal RNDs, for example, the surveys of Jackwerth (2004) and Christoffersen, Jacobs, and Chang (2011).

A distinct rich literature backs out implied risk-neutral moments from option prices, rather than the complete RND. Carr and Wu (2009) study implied variance and combine it with realized variance to build the variance risk premium (VRP). Bollerslev, Tauchen, and Zhou (2009) develop a model in which the VRP and equity risk premium share a common component, thus explaining why the VRP works empirically very well in predicting market returns. Our point of departure is that variance risk is inherently symmetric, while we explicitly investigate asymmetric volatility, which is larger conditional on a negative first-period return. Indeed, our empirical results suggest the presence of asymmetric volatility risk. Moreover, the particular functional form of asymmetric volatility matters, with the Frank copula fitting the data better than Student's t or Gaussian copulas.

Skinzi and Refenes (2004) and Driessen, Maenhout, and Vilkov (2005) infer the implied correlation between stocks in the index. Buss, Schoenleber, and Vilkov (2018) show both theoretically and empirically that the risk premium on such implied correlation between stocks and the VRP complement each other in predicting market returns. The VRP measures the fear of high volatility, especially due to jumps (Bollerslev and Todorov, 2011); the risk premium on correlation between stocks measures the risk of losing diversification. In contrast, we are interested in the asymmetric volatility effect, which assesses how volatile second-period returns are conditional on first-period returns.

Chang, Christoffersen, and Jacobs (2012) suggest market skewness as an alternative related factor. We document in our empirical work that we can also predict returns, volatility, and risk-neutral quantities, while documenting that the AVIC is different from variance and skewness risk.

3. Data and Definition of Variables

The Chicago Board of Options Exchange (CBOE) launched options on VIX futures in 2006, so we start the data collection on that date and end in August 2014. To create our weekly sample, we first find all short-dated (one-month) S&P 500 index option expiration dates. We also collect the matching short-dated (monthly) VIX futures option expiration dates and the long-dated (two-month) S&P 500 index option expiration dates, which expire one month later.[9] We then create four observation dates by going back two, three, four, and five weeks from the expiration date of the short-dated S&P 500 option. This procedure gives us about 400 weekly observation dates; however, due to limited data quality for VIX options, we drop some early data and select all dates from July 2007 to August 2014, which gives us 345 weekly observation dates in total.

[9]During our sample, short-dated S&P 500 options expire on the Saturday following the third Friday of a particular month. Long-dated S&P options expire similarly but one month later. VIX futures options expire on the Wednesday that is 30 days prior to the third Friday of the calendar month immediately following the expiring month.

3.1 Data on Index Returns

We obtain end-of-day underlying and options data on the S&P 500 index from OptionMetrics.[10] We work with midpoint implied volatilities inferred from raw option prices. We also record the S&P 500 dividend yields and interpolate the certificate of deposit rates from OptionMetrics to match the exact days-to-maturity of our S&P 500 options.

On each observation date, we collect out-of-the-money short-dated and long-dated S&P 500 index options, in which four related observation dates share the same expiration dates as detailed above. We eliminate options with zero bids and filter for moneyness (=strike price/index futures level) to lie between 1 ± 5 at-the-money implied volatilities observed on a given day and adjusted for the maturity of the options; see Andersen, Bondarenko, and Gonzalez-Perez (2015).[11] We normalize all option strikes by the index price. In order to compute RNDs and risk-neutral moments, we require at least eight index option prices to exist at any given observation date. We compute model-free values for implied variance and implied skewness following Bakshi, Kapadia, and Madan (2003). Following Bollerslev, Tauchen, and Zhou (2009), we compute the variance risk premium for the S&P 500 index as the difference between model-free implied variance for the next month and realized variance over the past month computed from high-frequency data.

3.2 Data on Future Expected Volatility

We obtain end-of-day data on options on VIX futures from OptionMetrics and daily VIX futures data from the CBOE. We collect the underlying VIX and VIX futures as daily closing data from the CBOE.[12] We work with midpoint implied volatilities inferred from VIX options. We estimate VIX implied volatilities from the Black (1976a) model using raw option prices and the reported VIX futures level at the end of the day.

[10] We use the IvyDB OptionMetrics database available through WRDS, updated in December 2014 to August 2014.

[11] We thank a referee for suggesting the use of the effective strike range.

[12] Gonzalez-Perez and Guerrero (2013) find intra-week effects in the VIX of -0.23% on average. Perturbing our data by twice such magnitude does not affect our results at all.

We use the same observation dates as for the S&P 500 options and collect the short-dated options on VIX futures. We eliminate zero bids and filter for moneyness (=strike price/VIX futures level) to lie between 1 ± 5 at-the-money implied volatilities adjusted for the maturity of the options. We normalize all option strikes by the VIX futures level. In order to compute RNDs, we require at least five available options on VIX futures.

4. Methodology

We first describe how we obtain the realized correlation between index returns and future expected volatility, the AVRC. Next, we work out the marginal RNDs of index returns and future expected volatility separately. Finally, we connect the two marginal RNDs via a copula, which allows us to work out the AVIC.

4.1 Asymmetric Volatility Realized Correlation Between Return and Future Expected Volatility

To compute the time-varying AVRC between returns and future expected volatility for each observation date, we use a 360-day historical window. Within this window, we compute 30-day returns and 30-day relative changes in levels of VIX futures. These two time series have the same lead-lag structure as the AVIC, as the return is the current return and the change in VIX is the change in future expected volatility. The AVRC is then the correlation coefficient between the two time series.

4.2 Marginal Risk-neutral Distributions for Returns and Future Expected Volatility

On each observation date, we estimate the RNDs of the short-dated S&P 500 returns and of future expected volatility using the "fast and stable method" of Jackwerth (2004). Given a trade-off parameter, a closed-form solution exists for fitting the implied volatilities of observed options best, while, at the same time, delivering the smoothest implied volatility smile. The optimal implied volatilities can be translated directly into risk-neutral probabilities. The same paper argues that, given some low number of observed option prices, the exact choice of method matters little for obtaining RNDs. We ensure that the RND produces a mean index return equal to the monthly risk-free rate minus the dividend yield (i.e., the mean of the traded asset return under the risk-neutral measure). We fine-tune the volatility of the RND so that model option prices best match observed option prices. Finally, we record the volatility of the RND and scale the RND to have unit volatility. We denote this RND at time t by $\hat{Q}_t(m)$, in which each discrete density function is defined by a number of return values $m_i,\ i = 1, \ldots, N$ and their associated probabilities $\hat{Q}_t(m)$.

For the RND of future expected volatility, we use short-dated options on VIX futures. We basically follow the same method as above. Because VIX is not traded and we use futures on VIX directly, we set both the interest rate and the dividend yield to zero. We denote the RND of future expected volatility at time t by $\hat{Q}_t^\sigma(m^\sigma)$, in which each discrete density function is defined by a number of future expected volatility values $m_i^\sigma,\ i = 1, \ldots, N^\sigma$ and their associated probabilities $\hat{Q}_t^\sigma(m^\sigma)$. For future expected volatility, the distribution is specified by the relative deviation in future conditional volatility from its expected level, and, hence, it automatically takes on a mean of one. Also, the RND of future expected volatility is not rescaled to unit volatility.

4.3 Joint Risk-neutral Distribution of First-period Returns and Future Expected Volatilities

By now, we have the two marginal RNDs in place, but not yet the bivariate distribution. We rely on Sklar's (1959) theorem that every multivariate cumulative distribution can be represented as a copula defined over the marginal distributions. We consider five different copulas: two implicit ones (Gaussian and Student's t) and three explicit ones (Frank, Clayton, and Gumbel). The description and formulas of each copula can be found in McNeil, Frey, and Embrechts (2005) and Trivedi and Zimmer (2007).

We select the best copula for fitting the bivariate distribution of index returns and expected future volatility under the physical probability measure. Here, we use historical time series of monthly S&P 500 index returns and monthly relative changes in VIX. Note that even though the changes are both computed over the same month, the returns are backward-looking, while relative changes in VIX are forward-looking, because they are changes to future *expected* volatility.

Overall, symmetric copulas (Frank, Student's t, and Gaussian) work well under the physical probability measure, while asymmetric ones (Clayton and Gumbel) do not. For each of the time series, we use kernel smoothing based on a Normal kernel with optimal bandwidth (Bowman and Azzalini, 1997) to estimate the cumulative density function and then estimate the respective copula parameters using maximum likelihood.[13] This gives us the unconditional dependency parameter for each copula θ, and for Student's t copula the degrees of freedom parameter ν. To test how well the dependency in real data is captured by each copula, we simulate 5,000 realizations from each copula and use the two-dimensional Kolmogorov-Smirnov test developed in Peacock (1983) to test the null that the simulated distribution is the same as the empirical one.

The results are collected in Table 1. For the empirical distribution of index returns and relative changes in VIX, the unconditional AVRC over the whole sample is -0.736. Note that this differs somewhat from our above average value of the (conditional) AVRC, which is -0.773. The difference is due to the use of overlapping moving windows in the 360-day historical samples used for the AVRC. Looking at the simulated AVRC for the five copulas, we find that they are all in the vicinity (-0.633 to -0.752) of the unconditional AVRC, but for the Clayton copula with a simulated AVRC of -0.020. Testing to see if the simulated distributions from the copulas are indistinguishable from the empirical distribution, we reject the test for the Clayton and Gumbel copulas with p-values of less than 0.01. We are thus left with the Frank, Student's t, and Gaussian copulas for possible use to model an AVIC.

[13] For the Gumbel copula we use the time series of index returns and the series of *minus* relative changes in VIX, because this copula is used to model positive dependency, and the expected direction of dependency in our case is negative.

Table 1: Copula calibrations under the physical measure

The table shows Frank, Student's t, Gaussian, Clayton, Gumbel copulas fitted to monthly index returns and monthly relative changes in future expected volatility (VIX) by maximum likelihood. The sample period is July 2007 to August 2014. The AVRC is based on the unconditional empirical distribution and (for the copulas) on a simulated distribution based on 5,000 draws from the copula. The column "KS test (p-value)" gives the p-value of the two-sided two-dimensional Komogorov-Smirnov test with the null that the simulated distribution is the same as the empirical one.

Model	AVRC	KS test (p-value)
Empirical distribution	-0.736	-
Frank	-0.648	0.245
Student's t	-0.752	0.199
Gaussian	-0.740	0.182
Clayton	-0.020	0.000
Gumbel	-0.633	0.009

The final decision on which copula to use cannot be made until we know which one performs best in terms of modeling a conditional AVIC. Skipping ahead to detailed comparisons in Section 6, this turns out to be the Frank copula, which we will thus use from now on.

The Frank (1979) copula is given by

$$C_\theta(u, u^\sigma) = -\frac{1}{\theta} \log \left(1 + \frac{(\exp(-\theta u) - 1)(\exp(-\theta u^\sigma)) - 1)}{\exp(-\theta) - 1} \right), \tag{1}$$

where θ is the dependence parameter. u and u^σ are random vectors with uniformly distributed marginals, where each realization of the random vector can be converted into the realization of the respective marginal (the RND of return $\hat{Q}_t$ and the RND of future expected volatility $\hat{Q}_t^\sigma$) by using the inverse cumulative marginal distribution. The copula produces a bivariate RND $\hat{Q}_t(r_{t,t_1}, \sigma_{t_1,t_2})$ of first-period (which can be two, three, four, or five weeks) index returns r_{t,t_1} and future expected volatility σ_{t_1,t_2}.

4.4 Risk-neutral Two-period Return Distribution and Asymmetric Volatility Implied Correlation

But how do we calibrate the parameter θ of the Frank copula? The idea is to use the bivariate RND of returns and future expected volatility to build up the *long-dated*, two-period index RND. We can then change θ in such a way as to minimize the pricing error of the long-dated S&P 500 options.

Given θ, we draw 200,000 pairs of uniformly distributed numbers $u,\ u^{\sigma} \sim U[0,1]$, where u is iid, and u^{σ} is drawn conditionally on u and the copula parameter θ.[14] The corresponding realizations of an index return and its future expected volatility are found by plugging u and u^{σ} into the respective inverse cumulative marginal distributions. Their joint distribution has the short-dated option-implied RNDs as marginals, and the Frank copula with parameter θ determines the bivariate RND.

Based on those realizations, we construct a *conditional* distribution of second-period index returns $\hat{Q}_t(r_{t_1,t_2}|r_{t,t_1},\sigma_{t_1,t_2})$ by starting with the distribution of normalized index returns $m \sim \hat{Q}_t$. Crucially, we assume that the conditional second-period returns are distributed as the first-period returns but with different volatility, drawn according to the RND of future expected volatility. We adjust the volatility of returns to match the drawn instance of second-period volatility σ_{t_1,t_2}, and, at that time, we take care that we scale mean and volatility to account for the number of days between t_1 and t_2. The resulting second-period return is then

$$r_{t_1,t_2} = (m - (r_f - \delta))\sigma_{t_1,t_2}\sqrt{t_2 - t_1} + (r_f - \delta)(t_2 - t_1), \tag{2}$$

where r_f is the annualized risk-free rate and δ is the annualized dividend yield. At this point we have $200,000$ second-period returns, each generated by a realization of first-period return and second-period future expected volatility.

[14]The numerical procedure converges after approximately 50,000 draws.

Having the second-period returns in place, we create the RND of two-period returns $\hat{Q}_t(r_{t,t_2})$. For each second-period return, we know its preceding first-period return. Aggregating the two returns gives us 200,000 two-period index returns, which we resample onto a return space of 500 equally spaced returns to obtain a two-period return RND.

We then use the two-period RND to price the long-dated index options. Each set of prices depends on two parameters, namely, the dependency parameter θ of the copula and the mean of the future volatility distribution $\bar{\sigma}_{t_1,t_2}$ that we use to rescale the volatility distribution.[15] We find the optimal parameters by minimizing the root-mean squared error of the relative deviation between model and market implied volatilities.[16] As the dependence parameter θ is not directly comparable to the AVRC computed above, we compute the AVIC based on the optimal θ as the correlation between the resulting 200,000 returns and expected future volatilities.[17]

For each observation date t, we define the *ex ante* AVCRP as the difference between the AVRC based on the past 360 days (from t^- to t) and the currently observed AVIC (based on the time to expiration of the short-dated option from t to t_1:

$$AVCRP_t = AVRC_{t^-,t} - AVIC_{t,t_1}. \tag{3}$$

By using non-matching periods in formula (3), we follow the definition of the *ex ante* VRP (as in Bollerslev, Tauchen, and Zhou, 2009) to avoid a forward-looking bias in the empirical analysis. The use of the past AVRC is justified by the fact that the AVRC is extremely persistent; we fail to reject a unit root using the Augmented Dickey-Fuller test with a p-value of 0.54.

---

[15] Note that we model the RND of relative changes of VIX futures from their expected value. However, there is no easily predetermined value for this expected value since volatility is not an easily traded asset. Thus, we need to pin it down by calibrating it jointly with the dependency parameter θ to match the long-dated index option prices.

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

[17] We compute the AVIC between returns and future expected volatility at the monthly frequency, ignoring the exact stochastic processes governing what happens in between months. We thank a referee for pointing out that the AVIC would only be the same as the conditional expected correlation under the risk-neutral measure for stationary underlying processes of returns and volatility. Such stationarity is not given in the data, as volatility is highly persistent and time-varying.

5. Empirical Results

We present results on the AVIC, which we interpret as a measure of asymmetric volatility risk. We relate it to other economically relevant variables and document how it predicts future returns, volatility, and risk-neutral implied quantities. In the subsequent analysis we will use the Frank copula and relegate results for the Gaussian and Student's t copulas to the robustness section.

5.1 Asymmetric Volatility Implied Correlation and Asymmetric Volatility Correlation Risk Premium

Table 2: Asymmetric volatility implied and realized correlation statistics

The table shows the mean, standard deviation, skewness, kurtosis, minimum, and maximum for asymmetric volatility implied correlation and asymmetric volatility realized correlation during our sample from July 2007 to August 2014. We also show the statistics for the pre-crisis period (July 2007 to August 2008), the crisis period (August 2008 to December 2009), and the post-crisis period (January 2010 to August 2014).

	Full period		Pre-crisis		Crisis period		Post-crisis	
	AVIC	AVRC	AVIC	AVRC	AVIC	AVRC	AVIC	AVRC
Mean	−0.833	−0.773	−0.863	−0.853	−0.892	−0.764	−0.807	−0.756
Standard dev.	0.137	0.095	0.126	0.042	0.071	0.104	0.148	0.092
Skewness	1.616	0.678	1.558	0.161	1.817	0.479	1.403	0.585
Kurtosis	5.891	2.576	4.851	1.743	7.176	2.063	5.072	2.469
Min	−0.987	−0.913	−0.987	−0.913	−0.982	−0.906	−0.987	−0.889
Max	−0.189	−0.488	−0.440	−0.779	−0.607	−0.568	−0.189	−0.488

Table 2, column Full period, shows descriptive statistics for the AVIC and the AVRC for the whole sample from July 2007 to August 2014. The AVRC has a mean of −0.77 with a standard deviation of 0.10, while AVIC has a mean of −0.83 with a standard deviation of 0.14. The AVCRP has a mean of 0.06 with a standard deviation of 0.17 and is strongly significant with a p-value of less than 0.01, based on a standard error of the mean (0.0150) adjusted with four lags, see Newey and West (1987). Before the financial crisis started in August 2008, the AVCRP was 0.01, which suggests that asymmetric volatility risk was barely priced before the crisis. The AVCRP increased to 0.13 during the crisis (August 2008 to December 2009). During this time, option prices implied a more negative AVIC (by −0.03), as future expected volatility is assumed to react stronger to negative returns than in calm times, even accounting for overall higher future expected volatility. Yet the AVRC is much less negative (by +0.09) because volatility decreases even after negative returns as the crisis abates. After the crisis (January 2010 to August 2014), the AVCRP contracts to 0.05 which the AVIC increases by 0.08 as the option markets relax.

Standard deviations are fairly large and contract somewhat for the AVIC during the crisis due the limited range of correlation. In crisis times, the AVIC approaches −1 from above but cannot go any lower, thus lowering standard deviations. The distributions of the AVIC and the AVRC are mildly positively skewed (1.62 and 0.68), and the AVIC is somewhat leptokurtic (5.89 vs. 2.58 for AVRC). Minimal values approach −1, and maximal values never exceed 0.

Depicting the time series of the AVIC and the AVRC, we see in Figure 2 that both are highly time-varying. The correlation between the two time series is only 0.06 and insignificant (p-value 0.25),[18] which indicates that they contain different information: the AVRC typically does not account for the rarely observed tail regions of the return distribution, while the AVIC is based on expectations over the whole range of returns, including the tails. Both series are highly persistent, and we fail to reject a unit root using the Augmented Dickey-Fuller test for the AVRC (p-value of 0.54) and the AVIC (p-value of 0.11).

5.2 Asymmetric Volatility Implied Correlation Versus Other Variables

Economically, a significant AVCRP means that taking on asymmetric volatility risk is being compensated in the market. Importantly, AVCRP is based on the joint distribution of returns and future expected volatility, whereas alternative risk measures (e.g., implied

[18] The cross correlation function does not show significant values from minus five to plus five lags.

Figure 2: Time series of asymmetric volatility implied and realized correlation

The figure shows weekly time series of asymmetric volatility implied (triangles) and realized (dots) correlation over the sample period from July 2007 to August 2014.

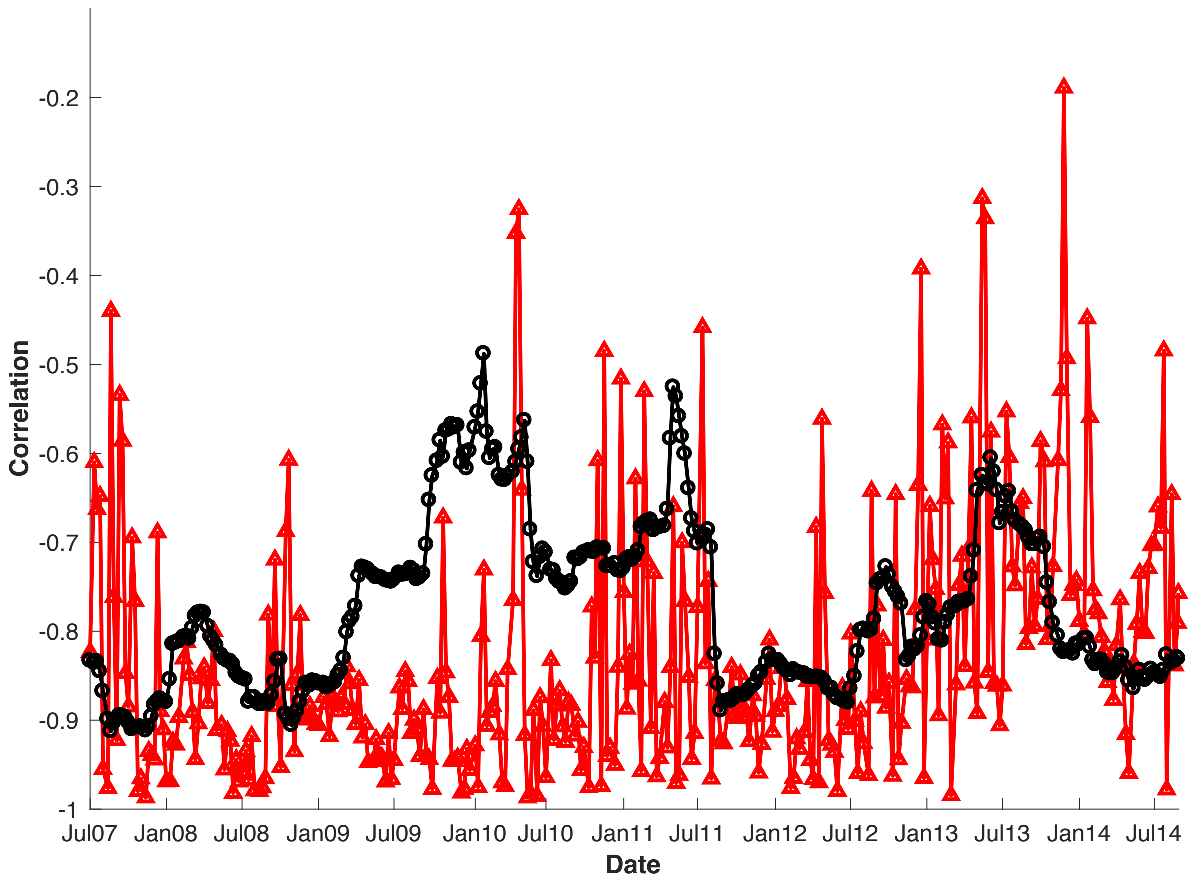

volatility, implied skewness, and VRP) are based only on the (risk-neutral and physical) return distribution.

To further explore, we compare the AVIC and the AVCRP to various variables that could be affected by the interplay of index returns and future expected volatility. We report correlations across seven variables in Table 3, Panel A, and the confidence bounds of those correlations in Panel B. The AVIC and the AVCRP are highly negatively correlated (-0.81), and in the following we only discuss the AVIC.

The VRP is the payment for a hedge against negative returns by gaining on increased volatility. Inherently, the VRP is symmetric as an up or down return (of same size and

Table 3: Correlations

The table shows in Panel A the correlations between the asymmetric volatility implied correlation, the asymmetric volatility correlation risk premium, the variance risk premium, VIX, long-dated minus short-dated option implied skewness, short-dated option implied skewness, and the realized S&P 500 return from t to t_1 (index return). The sample period is July 2007 to August 2014. All variables except the index returns are known on the observation date. Panel B shows the 5% confidence intervals. The lower bounds are in the lower triangular part of the table and the upper bounds in the upper part of the table.

Panel A: Sample Correlations

Variable	#	1	2	3	4	5	6	7
AVIC	1	1.00	-	-	-	-	-	-
AVCRP	2	−0.81	1.00	-	-	-	-	-
VRP	3	0.19	−0.05	1.00	-	-	-	-
VIX	4	−0.25	0.07	−0.29	1.00	-	-	-
Implied Skewness, long-short	5	0.32	−0.20	0.21	−0.37	1.00	-	-
Implied Skewness, short	6	−0.34	0.21	−0.32	0.57	−0.89	1.00	-
Index return, first-period	7	0.03	0.06	0.08	−0.02	−0.05	0.00	1.00

Panel B: Confidence Bounds for Sample Correlations

Variable	#	1	2	3	4	5	6	7
AVIC	1	-	−0.77	0.29	−0.15	0.41	−0.24	0.14
AVCRP	2	−0.84	-	0.06	0.18	−0.09	0.31	0.16
VRP	3	0.09	−0.16	-	−0.19	0.31	−0.22	0.19
VIX	4	−0.34	−0.03	−0.38	-	−0.28	0.64	0.09
Implied Skewness, long-short	5	0.22	−0.30	0.11	−0.46	-	−0.87	0.05
Implied Skewness, short	6	−0.43	0.11	−0.41	0.50	−0.91	-	0.10
Index return, first-period	7	−0.07	−0.05	−0.02	−0.12	−0.16	−0.11	-

relative to the mean) results in the same contribution to variance.[19] A first indication that the AVIC measures something different from the symmetric VRP can be found in the moderate correlation with the VRP of 0.19. As a second symmetric risk measure, we consider VIX, which is often referred to as a market fear gauge. It is negatively correlated with the VRP (−0.29) and the AVIC (−0.25).

The AVIC closer to its lower limit of −1 affects the two-period return RND through the left tail in that a negative first-period return induces high volatility during the second

[19]Bollerslev and Todorov (2011) claim that the VRP is mostly driven by jumps and not the diffusive part of the returns. However, their measures of jumps for the left tail and the right tail are highly correlated. Our argument that the VRP is mainly about symmetric tail risks thus remains valid.

period and thus enlarges the left tail. The effect is weaker for the right tail. We thus include as variables the implied skewness from short-dated options and implied skewness from long-dated options minus implied skewness from short-dated options. Both skewness measures are correlated with the AVIC with absolute correlations of up to 0.34, and they are also highly correlated with each other (-0.89). We concentrate on long-dated minus short-dated implied skewness henceforth, as it more closely mirrors the intuition of the AVIC in relating first-period to two-period skewness.

We finally observe that the AVIC is positively but insignificantly correlated with the first-period index return (0.03). This suggests that the AVIC might be capable of predicting index returns; we next investigate this exciting proposition further. Since we find the AVIC to be economically related to VIX, the VRP, and long-dated minus short-dated implied skewness, although correlations tend to be moderate, we will use those variables as controls in further explorations.

5.3 Predicting S&P 500 Returns and Volatilities

In our regression analysis, we predict returns, volatility, and an *ex post* VRP (from t to the maturity of the short-dated options) with the AVIC at time t while controlling for VIX, the (*ex ante*) VRP, and long-dated minus short-dated implied skewness.

We see in Table 4 that a more negative AVIC is associated with higher future returns. The results are insignificant (p-value of 0.14) without controls but strengthen when controlling for VIX (p-value of 0.01) or the VRP (p-value of 0.02). When using controls, we always include long-dated minus short-dated implied skewness.

The results for volatility go into the opposite direction when we control for VIX (p-value of 0.00). This is due to the strong negative correlation of volatility with returns. Results are insignificant in the regression without controls and the regression with the VRP as a control.

Table 4: Predicting S&P 500 returns, volatility, and the *ex post* variance risk premium

We define our dependent variables as realized return, volatility, and the *ex post* VRP, measured from t to the maturity of the short-dated options. Independent variables are the asymmetric volatility implied correlation, VIX, the variance risk premium, and long-dated minus short-dated implied skewness. The p-values are based on the Newey and West (1987) standard errors. The sample period is July 2007 to August 2014.

	Return			Volatility			*Ex post* VRP		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Intercept	−0.359	−0.378	−0.752	0.109	0.080	0.295	1.255	0.820	0.464
	0.205	0.254	0.014	0.006	0.006	0.000	0.000	0.052	0.537
AVIC	−0.489	−0.732	−0.706	−0.084	0.103	0.035	−0.916	−1.438	−1.458
	0.136	0.013	0.020	0.104	0.002	0.909	0.040	0.002	0.002
VIX	-	−0.915	-	-	0.855	-	-	−0.591	-
	-	0.322	-	-	0.000	-	-	0.327	-
VRP	-	-	0.095	-	-	−0.032	-	-	0.109
	-	-	0.168	-	-	0.000	-	-	0.177
Implied Skewness, long-short	-	0.065	0.107	-	−0.023	−0.078	-	0.377	0.391
	-	0.610	0.138	-	0.000	0.000	-	0.003	0.001
Adjusted R^2	0.005	0.019	0.014	0.007	0.591	0.146	0.013	0.047	0.051

As the VRP is the difference of implied and realized volatility, we can also predict the *ex post* VRP with p-values of less than 0.04 for all of our three models.

Generally, results for returns, volatility, and *ex post* VRP are even stronger if we use the AVCRP instead of the AVIC.

5.4 Asymmetric Volatility Implied Correlation Compared with Implied Risk Measures and Moments

We now investigate how second-period (i.e., observed at time t_1 and inferred from the second-period RND from t_1 to t_2) risk-neutral quantities react to changes in the AVIC. For our tail measure "Right minus left tail", we look at the difference between the probability mass in the right tail and the left tail of the distribution, divided by the probability between the tails. We use a $\pm 10\%$ return per month threshold for identifying the tail regions, which corresponds to about $\pm 2\sigma$ (using an unconditional index volatility σ of 0.17 p.a.). Implied skewness and implied volatility are averages computed from the second-period simulated RNDs. Again, we control for VIX, the VRP, and long-dated minus short-dated implied skewness. We use non-overlapping regressions to guarantee that the quantity we want to predict is always observed about one month after the estimation date t.

Table 5: Predicting future risk-neutral quantities for the S&P 500

We use as dependent variables the selected quantities of the future option-implied distribution of S&P 500 estimated on the expiration date t_1 of the short-dated options that are used to calibrate the AVIC at time t. "Right minus left tail" is the probability mass to the right of the 10% monthly return minus the probability mass to the left of the −10% monthly return (which corresponds approximately to the $\pm 2\sigma$ event using an unconditional index volatility σ of 0.17 p.a.), divided by the probability between the tails. Implied skewness and implied volatility are computed from the second-period simulated return RNDs. Independent variables are the asymmetric volatility implied correlation, VIX, the variance risk premium, and long-dated minus short-dated implied skewness. We use non-overlapping regressions to guarantee that the quantity we want to predict is always observed about one month after the estimation date t. For p-values we use White (1980) standard errors adjusted for heteroskedasticity. The sample period is July 2007 to August 2014.

	Right *minus* left tail			Implied skewness			Implied volatility		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Intercept	0.004	−0.025	−0.005	−2.135	−1.740	−1.167	0.045	0.028	0.084
	0.404	0.413	0.373	0.000	0.000	0.000	0.034	0.107	0.000
AVIC	0.053	0.043	0.039	−0.801	0.026	−0.069	−0.024	0.012	0.003
	0.020	0.087	0.142	0.038	0.167	0.366	0.581	0.893	0.184
VIX	-	0.067	-	-	2.058	-	-	0.212	-
	-	0.038	-	-	0.000	-	-	0.000	-
VRP	-	-	−0.004	-	-	−0.077	-	-	−0.005
	-	-	0.837	-	-	0.461	-	-	0.581
Implied Skewness, long-short	-	0.023	0.018	-	−0.647	−0.832	-	−0.007	−0.028
	-	0.063	0.213	-	0.000	0.000	-	0.729	0.033
Adjusted R^2	0.065	0.126	0.073	0.052	0.556	0.363	0.002	0.600	0.095

Our results in Table 5 suggest that a more negative AVIC, that is, a higher asymmetric volatility risk, predicts a fatter left tail of the index RND observed in the future (at t_1) (p-value of 0.02). Controlling for VIX increases the p-value to 0.09; controlling for the VRP turns the regression insignificant (p-value of 0.14).

A more negative AVIC predicts an increase in future implied skewness (p-value of 0.04). This significance is lost, though, once we add controls. The relation is related to the fattening of the left tail as the AVIC goes more negative. Such fattening on the left needs to be balanced by more probability mass to the right of the mean, while the risk-neutral mean is fixed by the risk-free rate and the dividend yield. These movements reduce the typical left-skewness by making the distribution more symmetric. The AVIC does not predict second-period implied volatility.

The results show that the AVIC predicts asymmetric tail behavior of the future (second-period) index RND, which contains important information about future investment opportunities. The use of the AVIC, which incorporates information about both returns and future expected volatility, seems to offer benefits, compared to the use of information on only one or the other quantity.

5.5 Bivariate RND of Returns and Future Expected Volatilities

As we have estimates in hand of the bivariate RND of returns and future expected volatility, we can price basket options written on a combination of the two quantities. The AVIC could then inform the risk management of banks holding assets simultaneously linked to index returns and volatility. Furthermore, the time-varying nature of our findings suggests a time-varying and potentially priced role for asymmetric volatility risk. This is relevant for the development of stochastic volatility option pricing models, which mostly model the correlation between returns and contemporaneous volatility as a constant and cannot consider a price of risk for that contemporaneous correlation.

Our method can also be used to compute term structures of the AVIC. Here we compute the AVIC not of first-period returns and second-period future expected volatility but for t-period returns and $t+1$-period future expected volatility. We find those RNDs from options on returns that expire in t periods and options on VIX 30-day futures that also expire in t periods.

Areas of exciting development are markets other than the S&P 500 index, for which the AVIC can be computed. On the CBOE, options on the underlying and VIX-type contracts exists for seven stock indices, five individual stocks, 10 commodity and country ETFs, two interest rates, four currencies, and five volatilities. Interestingly, CBOE filed in 2011 with the SEC for permission to introduce options on 40 VIX contracts, many of them from the list above.[20] As of 2018, no such options on VIX are traded yet, but once they are introduced, our methodology would also apply to all these new markets.

6. Robustness

In Section 4.3, we looked at the performance of various copulas under the physical measure. Now we are ready to also compare performance under the risk-neutral measure. We replace the Frank copula used above by one of the alternative copulas and repeat the AVIC optimization. We report the pricing errors (defined as the average root mean squared relative deviation of model-based implied volatilities from observed ones, always using the optimal copula parameter θ) in Table 1. We see that the Frank copula delivers the lowest pricing error (0.0332) among all copulas. The Gaussian (0.0355) and Student's t (0.0336) pricing errors are close. The Clayton (0.0344) pricing error is somewhat further away, but here the AVIC is shifted much lower (-0.94 compared to -0.83 for Frank). The Gumbel copula does not work well at all, with a pricing error of 0.1031 and a positive AVIC.

Table 6: Copula calibrations under the risk-neutral measure

The table shows the average (across all observation dates) asymmetric volatility implied correlation from the simulated distributions and pricing errors computed as the average root mean squared relative deviation of model-based implied volatilities from observed ones for the Frank, Student's t, Gaussian, Clayton, Gumbel copulas. The sample period is July 2007 to August 2014.

Model	AVIC	Pricing Errors
Frank	−0.8328	0.0332
Student's t	−0.8478	0.0336
Gaussian	−0.8350	0.0335
Clayton	−0.9432	0.0344
Gumbel	0.1054	0.1031

We further investigate how sensitive the three better performing copulas (Frank, Student's t, and Gaussian) perform to variation in the parameter θ. For each date, we compute the AVIC and pricing error (average root mean squared relative deviation of

[20] http://ir.cboe.com/press-releases/2011/16-mar-2011c.aspx

model-based implied volatilities from observed ones) associated with different parameters θ for our three copulas. We then compute relative deviations (in %) of pricing error from the minimal pricing error for a range $(-0.1, 0.1)$ of relative deviations of the AVIC from the optimal AVIC, in which the optimal AVIC and the corresponding minimal pricing error are based on the Frank copula. We average those relative deviations from the minimal error along the timeline and depict them in Figure 3. The lowest pricing error is reached for the Frank copula at the optimal AVIC. Deviating from the optimal AVIC by around 6% on either side increases the pricing error by 1%, compared to the optimal pricing error. Performances of Student's t and the Gaussian copulas are worse, with pricing errors being more than 3% higher. The sensitivity of those copulas to deviations from the optimal AVIC is similar to the sensitivity of the Frank copula. We thus feel vindicated to use the Frank copula in our main results.

Yet for robustness' sake, we repeat our main results on predicting physical and risk-neutral quantities for Student's t and Gaussian copulas.

Using the two stronger competing copulas (Student's t and Gaussian) to predict S&P 500 returns, volatilities, and the *ex post* VRP, results in Table 7 weaken when compared to the main results in Table 4 based on the Frank copula. For both alternative copulas, the AVIC still predicts realized volatility in the regression without controls, but the other coefficients turn insignificant. It emerges that the choice of copula matters and the Frank copula exhibits lower pricing errors and can better relate the AVIC to economically relevant quantities than Student's t or Gaussian copulas.

Predicting risk-neutral quantities for Student's t and Gaussian copulas generates results in Table 8 that are similar to our main results in Table 5 based on the Frank copula. All previously significant results remain so with the coefficient for the AVIC, now showing p-values of 0.05 (up from 0.02) for univariately predicting "Right *minus* left tail" and the coefficient for the AVIC, now showing p-values of 0.01 (down from 0.04) for predicting implied skewness. The results generally confirm our initial intuition that the Frank copula is the best choice for our main results.

Figure 3: Average pricing errors

The figure shows the (time-averaged) relative deviation (in %) of pricing errors (average root mean squared relative deviation of model-based implied volatilities from observed ones) for the Frank (solid line), Gaussian (dot-dashed line), and Student's t (dashed line) copulas with respect to pricing error based on the optimal asymmetric volatility implied correlation of the Frank copula. On the x-axis, we vary the AVIC of the three copulas by varying the parameter θ and record the relative deviation from the associated AVIC from the optimal AVIC under the Frank copula.

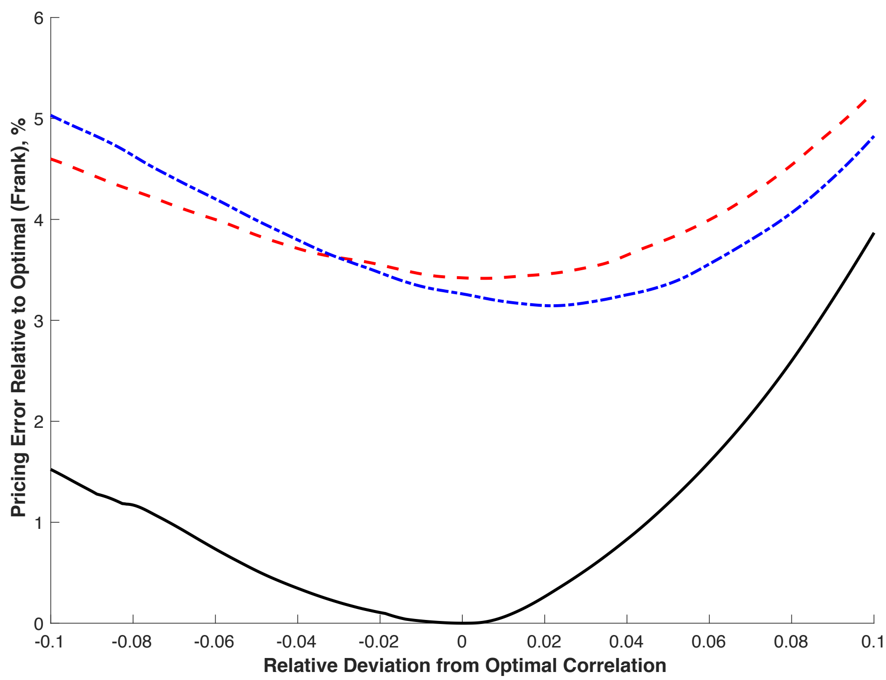

Table 7: Predicting S&P 500 returns, volatility, and the *ex post* variance risk premium

We define our dependent variables as realized return, volatility, and the *ex post* VRP, measured from the observation date to the expiration date of the nearest maturity options. Independent variables are the asymmetric volatility implied correlation, VIX, the variance risk premium, and long-dated minus short-dated implied skewness. The p-values are based on the Newey and West (1987) standard errors. The sample period is July 2007 the August 2014. In Panel A we use Student's t copula and in Panel B the Gaussian copula.

Panel A: Student's t copula

	Return			Volatility			*Ex post* VRP		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Intercept	−0.005	0.057	−0.304	0.089	0.030	0.232	1.675	1.506	1.140
	0.034	0.388	0.452	0.007	0.408	0.000	0.000	0.000	0.003
AVIC	−0.062	−0.207	−0.202	−0.106	0.043	−0.033	−0.405	−0.622	−0.682
	0.382	0.793	0.832	0.008	0.225	0.741	0.467	0.153	0.109
VIX	-	−0.838	-	-	0.849	-	-	−0.508	-
	-	0.402	-	-	0.000	-	-	0.474	-
VRP	-	-	0.089	-	-	−0.030	-	-	0.108
	-	-	0.217	-	-	0.001	-	-	0.196
Implied Skewness, long-short	-	0.013	0.053	-	−0.016	−0.074	-	0.280	0.289
	-	0.330	0.700	-	0.008	0.000	-	0.029	0.017
Adjusted R^2	−0.003	0.005	0.001	0.016	0.581	0.147	0.001	0.020	0.024

Panel B: Gaussian copula

	Return			Volatility			*Ex post* VRP		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Intercept	0.004	0.065	−0.288	0.087	0.027	0.229	1.667	1.497	1.143
	0.025	0.413	0.528	0.017	0.550	0.000	0.000	0.000	0.006
AVIC	−0.053	−0.200	−0.190	−0.110	0.040	−0.037	−0.420	−0.640	−0.695
	0.301	0.912	0.964	0.013	0.341	0.717	0.511	0.184	0.145
VIX	-	−0.831	-	-	0.847	-	-	−0.500	-
	-	0.407	-	-	0.000	-	-	0.482	-
VRP	-	-	0.087	-	-	−0.030	-	-	0.105
	-	-	0.222	-	-	0.001	-	-	0.217
Implied Skewness, long-short	-	0.013	0.052	-	−0.016	−0.074	-	0.280	0.289
	-	0.320	0.711	-	0.008	0.000	-	0.030	0.017
Adjusted R^2	−0.003	0.005	0.001	0.016	0.581	0.147	0.001	0.020	0.024

Table 8: Predicting future risk-neutral quantities for the S&P 500

We use as dependent variables the selected quantities of the future option-implied distribution of S&P 500 estimated *on the expiration date* t_1 of the short-dated options that are used to calibrate the AVIC at time t. "Right minus left tail" is the probability mass to the right of the 10% monthly return minus the probability mass to the left of the -10% monthly return (which corresponds approximately to the $\pm 2\sigma$ event using an unconditional index volatility σ of 0.17 p.a.), divided by the probability between the tails. Implied skewness and implied volatility are computed from the second-period simulated return RNDs. Independent variables are the asymmetric volatility implied correlation, VIX, the variance risk premium, and long-dated minus short-dated implied skewness. We use non-overlapping regressions to guarantee that the quantity we want to predict is always observed about one month after the estimation date t. For p-values we use White (1980) standard errors adjusted for heteroskedasticity. The sample period is July 2007 to August 2014. In Panel A we use Student's t copula and in Panel B the Gaussian copula.

Panel A: Student's t copula

	Right *minus* left tail			Implied skewness			Implied volatility		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Intercept	−0.004	−0.032	−0.012	−2.184	−1.847	−1.351	0.032	0.017	0.064
	0.450	0.157	0.991	0.000	0.000	0.000	0.115	0.393	0.003
AVIC	0.042	0.036	0.030	−0.853	−0.102	−0.282	−0.039	−0.002	−0.020
	0.052	0.116	0.255	0.012	0.702	0.581	0.113	0.183	0.673
VIX	-	0.071	-	-	2.038	-	-	0.211	-
	-	0.028	-	-	0.000	-	-	0.000	-
VRP	-	-	−0.004	-	-	−0.079	-	-	−0.005
	-	-	0.777	-	-	0.434	-	-	0.535
Implied Skewness, long-short	-	0.025	0.020	-	−0.626	−0.793	-	−0.005	−0.024
	-	0.035	0.134	-	0.000	0.000	-	0.931	0.070
Adjusted R^2	0.046	0.120	0.062	0.075	0.557	0.372	0.031	0.597	0.105

Panel B: Gaussian copula

	Right *minus* left tail			Implied skewness			Implied volatility		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Intercept	−0.002	−0.029	−0.011	−2.287	−1.937	−1.445	0.031	0.017	0.064
	0.163	0.237	0.863	0.000	0.000	0.000	0.159	0.446	0.004
AVIC	0.046	0.040	0.033	−0.994	−0.213	−0.405	−0.040	−0.002	−0.021
	0.045	0.108	0.243	0.005	0.752	0.315	0.127	0.200	0.694
VIX	-	0.071	-	-	2.023	-	-	0.211	-
	-	0.029	-	-	0.000	-	-	0.000	-
VRP	-	-	−0.004	-	-	−0.083	-	-	−0.006
	-	-	0.843	-	-	0.385	-	-	0.509
Implied Skewness, long-short	-	0.024	0.020	-	−0.611	−0.773	-	−0.005	−0.024
	-	0.037	0.147	-	0.000	0.000	-	0.925	0.073
Adjusted R^2	0.049	0.122	0.063	0.091	0.560	0.378	0.029	0.597	0.105

7. Conclusion

In our exploration of asymmetric volatility risk, we began by measuring the asymmetric volatility realized correlation between S&P 500 index returns and (relative changes in) future realized volatility. Using a novel identification strategy, we also manage to work out its risk-neutral counterpart, the asymmetric volatility implied correlation, from short-dated options on VIX futures and short- and long-dated index options. The asymmetric volatility correlation risk premium (asymmetric volatility realized minus implied correlation) is positive and significantly different from zero.

A more negative asymmetric volatility implied correlation predicts a higher return and a higher *ex post* variance risk premium. It further predicts second-period risk-neutral quantities, namely a higher probability of a future market crash (tail probability). These results remain in place when we control for VIX, the *ex ante* variance risk premium, and long-dated minus short-dated implied skewness.

Our work provides new ideas for risk management and the pricing of portfolios of index and volatility related securities. Extensions are possible to term structures of asymmetric volatility implied correlation. Currently, permission to introduce options on VIX for 40 further markets has been requested from the SEC, widening the applications of our method. In the future, we intend to look at bivariate pricing kernels and asset pricing tests, which could use asymmetric volatility implied correlation as a factor.

References

Aït-Sahalia, Y., Fan, J., Laeven, R. J. A., Wang, C. D., and Yang, X. (2017) Estimation of the continuous and discontinuous leverage effects, *Journal of the American Statistical Association* 112, 1744–1758.

Aït-Sahalia, Y., Fan, J., and Li, Y. (2013) The leverage effect puzzle: Disentangling sources of bias at high frequency, *Journal of Financial Economics* 109, 224–249.

Amato, J. D., and Gyntelberg, J. (2005) CDS index tranches and the pricing of credit risk correlations, *BIS Quarterly Review* March, 73–87.

Andersen, T. G., Bondarenko, O., and Gonzalez-Perez, M. T. (2015) Exploring return dynamics via corridor implied volatility, *Review of Financial Studies* 28, 2902–2945.

Augustin, P., Subrahmanyam, M. G., Tang, D. Y., and Wang, S. Q. (2014) Credit default swaps: A survey, *Foundations and Trends in Finance* 9, 1–196.

Bakshi, G. S., Kapadia, N., and Madan, D. B. (2003) Stock return characteristics, skew laws, and the differential pricing of individual equity options, *Review of Financial Studies* 16, 101–143.

Bardgett, C., Gourier, E., and Leippold, M. (2017) Inferring volatility dynamics and risk premia from the S&P 500 and VIX markets, *Working Paper, University of Zurich*.

Bates, D. S. (2012) U.S. stock market crash risk, 1926–2010, *Journal of Financial Economics* 105, 229–259.

Bekaert, G., and Wu, G. (2000) Asymmetric volatility and risk in equity markets, *Review of Financial Studies* 13, 1–42.

Black, F. (1976a) The pricing of commodity contracts, *Journal of Financial Economics* 3, 167–179.

——— (1976b) Studies of stock price volatility changes, in *Proceedings of the 1976 Meetings of the American Statistical Association, Business and Economics Statistics Section* pp. 177–181.

Bollerslev, T., Tauchen, G., and Zhou, H. (2009) Expected stock returns and variance risk premia, *Review of Financial Studies* 22, 4463–4492.

Bollerslev, T., and Todorov, V. (2011) Tails, fears, and risk premia, *Journal of Finance* 66, 2165–2211.

Bowman, A. W., and Azzalini, A. (1997) *Applied Smoothing Techniques for Data Analysis: The Kernel Approach with S-Plus Illustrations*. Oxford University Press Oxford, England.

Branger, N., Kraftschik, A., and Völkert, C. (2014) The fine structure of variance: Consistent pricing of VIX derivatives, *Working paper, Münster University.*

Buraschi, A., Trojani, F., and Vedolin, A. (2014) When uncertainty blows in the orchard: Comovement and equilibrium volatility risk premia, *Journal of Finance* 69, 101–137.

Buss, A., Schoenleber, L., and Vilkov, G. (2018) Expected stock returns and the correlation risk premium, *Working paper, Frankfurt School of Finance and Management.*

Carr, P., and Madan, D. B. (2013) Joint modeling of VIX and SPX options at a single and common maturity with risk management applications, *Working Paper, New York University.*

Carr, P., and Wu, L. (2009) Variance risk premiums, *Review of Financial Studies* 22, 1311–1341.

Chang, B. Y., Christoffersen, P., and Jacobs, K. (2012) Market skewness risk and the cross-section of stock returns, *Journal of Financial Economics* 107, 46–68.

Christie, A. A. (1982) The stochastic behavior of common stock variances: Value, leverage and interest rate effects, *Journal of Financial Economics* 10, 407–432.

Christoffersen, P., Heston, S., and Jacobs, K. (2009) The shape and term structure of the index option smirk: Why multifactor stochastic volatility models work so well, *Management Science* 55, 1914–1932.

Christoffersen, P., Jacobs, K., and Chang, B. Y. (2011) Forecasting with option implied information, *CREATES Research Papers.*

Dennis, P., Mayhew, S., and Stivers, C. (2006) Stock returns, implied volatility innovations, and the asymmetric volatility phenomenon, *Journal of Financial and Quantitative Analysis* 41, 381–406.

Driessen, J., Maenhout, P., and Vilkov, G. (2005) Option-implied correlations and the price of correlation risk, *Working Paper, INSEAD.*

Duan, J.-C., and Yeh, C.-Y. (2012) Price and volatility dynamics implied by the VIX term structure, *Working paper, National University of Singapore (NUS), Business School and Risk Management Institute.*

Figlewski, S., and Wang, X. (2000) Is the 'leverage effect' a leverage effect?, *Working paper, New York University.*

Frank, M. (1979) On the simultaneous associativity of $f(x, y)$ and $x + yf(x, y)$, *Aequationes Mathematicae* 19, 194–226.

Fuertes, C., and Papanicolaou, A. (2014) Implied filtering densities on the hidden state of stochastic volatility, *Applied Mathematical Finance* 21, 483–522.

Glosten, L., Jagannathan, R., and Runkle, D. (1993) On the relation between the expected value and the volatility of nominal excess return on stocks, *Journal of Finance* 48, 1779–1801.

Gonzalez-Perez, M. T., and Guerrero, D. E. (2013) Day-of-the-week effect on the VIX. A parsimonious representation, *North American Journal of Economics and Finance* 25, 243–260.

Hasanhodzic, J., and Lo, A. (2011) Black's leverage effect is not due to leverage, *Working Paper, MIT.*

Jackwerth, J. (2004) Option-implied risk-neutral distributions and risk aversion, Research Foundation of AIMR, CFA Institute Charlottesville, Virginia, USA.

Kalnina, I., and Xiu, D. (2015) Nonparametric estimation of the leverage effect using information from derivatives markets, *Working paper, University of Chicago, Booth School of Business.*

McNeil, A. J., Frey, R., and Embrechts, P. (2005) *Quantitative Risk Management: Concepts, Techniques and Tools.* Princeton University Press Princeton.

Mueller, P., Stathopoulos, A., and Vedolin, A. (2017) International correlation risk, *Journal of Financial Economics* 126, 270–299.

Nelson, D. B. (1991) Conditional heteroskedasticity in asset returns: A new approach, *Econometrica* 59, 347–370.

Newey, W. K., and West, K. D. (1987) A simple, positive-semidefinite, heteroskedasticity and autocorrelation consistent covariance matrix, *Econometrica* 55, 703–708.

Peacock, J. A. (1983) Two-dimensional goodness-of-fit testing in astronomy, *Monthly Notices Royal Astronomy Society* 202, 615–627.

Skinzi, V. D., and Refenes, A.-P. N. (2004) Implied correlation index: A new measure of diversification, *Journal of Futures Markets* 25, 171–197.

Song, Z., and Xiu, D. (2016) A tale of two option markets: State-price densities and volatility risk, *Journal of Econometrics* 190, 176–196.

Trivedi, P., and Zimmer, D. (2007) *Copula Modeling: An Introduction for Practitioners* vol. 1 of *Foundations and Trends in Econometrics.* Now Publishers, Inc. Hanover, Massachusetts, USA.

Wang, C. D., and Mykland, P. A. (2014) The estimation of leverage effect with high-frequency data, *Journal of the American Statistical Association* 109, 197–215.

White, H. (1980) A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity, *Econometrica* 48, 817–838.