



The Hazards of Volatility Diversification

Carol Alexander

ICMA Centre, University of Reading

Dimitris Korovilas

ICMA Centre, University of Reading

First Version: January 2011

ICMA Centre Discussion Papers in Finance DP2011-04

Copyright © 2011 Alexander and Korovilas. All rights reserved.

ICMA Centre • The University of Reading
Whiteknights • PO Box 242 • Reading RG6 6BA • UK
Tel: +44 (0)1183 788239 • Fax: +44 (0)1189 314741
Web: www.icmacentre.rdg.ac.uk
Director: Professor John Board, Chair in Finance
The ICMA Centre is supported by the International Capital Market Association



The Hazards of Volatility Diversification

Carol Alexander and Dimitris Korovilas

This draft: February 4, 2011

Abstract

Recent research advocates volatility diversification for long equity investors. It can even be justified when short-term expected returns are highly negative, but only when its equilibrium return is ignored. Its advantages during stock market crises are clear but we show that the high transactions costs and negative carry and roll yield on volatility futures during normal periods would outweigh any benefits gained unless volatility trades are carefully timed. Our analysis highlights the difficulty of predicting when volatility diversification is optimal. Hence institutional investors should be sceptical of studies that extol its benefits. Volatility is better left to experienced traders such as speculators, vega hedgers and hedge funds.

JEL Classification: G11, G15, G23

Keywords: Black-Litterman Model; Institutional Investors; Mean-Variance Criterion; Optimal Asset Allocation; S&P 500; SPY ETF; VIX Futures

Authors Details:

Carol Alexander

Chair of Risk Management and Director of Research,
ICMA Centre, University of Reading,
Reading, RG6 6BA, UK.

Email: c.alexander@icmacentre.rdg.ac.uk

Dimitris Korovilas

Doctoral Candidate,
ICMA Centre, University of Reading,
Reading, RG6 6BA, UK.

Email: d.korovilas@icmacentre.ac.uk

Equity and commodities have become more highly correlated since the banking crisis. Between January 2005 and December 2007 the sample correlation between the daily returns on the S&P 500 stock index and those on the S&P Goldman Sachs Commodity index was only 0.02, but between January 2008 and December 2010 it rose to 0.43. In fact most traditional asset classes have become more highly correlated and this encourages equity investors to seek alternative means of diversification.

The large negative correlation between daily returns on the S&P 500 and those on the VIX volatility index, averaging about -0.7 before the banking crisis, became even more negative (-0.85) during the crisis. The correlation between equity and credit is also negative but it is not as great as the equity-volatility correlation. Moreover, volatility trading is surging whilst credit trading seems well past its peak of popularity. Thus it is volatility that is currently being extolled as the new, effective diversifier for traditional asset classes.

This article examines the pitfalls of volatility diversification by long equity investors. A survey of the literature on volatility trading focuses on papers that report benefits from using VIX futures and options or variance swaps for diversification. Then we describe the exchanged-traded volatility market and empirically investigate the trading characteristics of VIX futures, because we will examine when buy-and-hold positions on these may successfully diversify an S&P 500 exposure. Short volatility positions are not considered because they would be suboptimal for long equity investors, offering a positively correlated exposure with risk-return characteristics that are inferior to those of the equity index. We examine the assumptions of portfolio theory that underpin the benefits of volatility diversification and ask whether the timings of VIX futures trades are important, and if they can easily be predicted.

The existing literature is extended in the following ways: (i) we analyse the carry and rollover costs on buy-and-hold VIX futures positions based on different rollover methodologies; (ii) we give a clear message on the timing of volatility diversification by long equity investors; (iii) we compute the minimum expected return on volatility that will justify diversifying into VIX futures, under both Markowitz [1952] and Black and Litterman [1992] optimization, and examine whether such expectations are borne out in practice.

LITERATURE REVIEW

Several recent papers advocate the advantages of volatility as a useful diversifier for equity. Szado [2009] considers the diversification of S&P 500 exposure using VIX futures and options and SPX put options, claiming that a long volatility exposure is beneficial for diversification and that VIX derivatives are more efficient than an exposure through SPX options. But they only consider the crash period, the allocation between volatility and other assets is done arbitrarily, and they do not consider different rollover strategies. Lee and Lin [2010] do examine the optimal rollover strategy for volatility futures, but they consider volatility as a hedge rather than a diversifier. For hedging the S&P 500 its own futures are most effective and have lower transaction costs than VIX futures. Chen, Chung, and Ho [2010] test the in-sample diversification benefits to different equity portfolios by adding spot VIX, VIX futures and VIX-squared portfolios replicated by SPX call and put options. Optimal portfolio weights are positive for spot VIX and negative for VIX-squared portfolios. VIX futures offer a better portfolio in only two cases, and only if shorted. Their out-of-sample tests include only spot VIX and VIX-squared portfolios.

Several studies, e.g. Daigler and Rossi [2006], Dash and Moran [2005] and P ezier and White [2008] use the spot VIX, which is not tradeable. Moran and Dash [2007] show that the desirable qualities of VIX do not always carry over to VIX futures and options. Hafner and Wallmeier [2008] use an ex-post analysis to demonstrate the benefits of adding variance swaps to European equity portfolios. Egloff, Leippold, and Wu [2010] also focus on variance swaps, modelling their dynamics in a two-factor model and promoting their diversification benefits as well as those of volatility futures.

The Black and Litterman [1992] model is used in our ex-ante analysis. It combines equilibrium returns with expected returns (in the form of investor’s current beliefs or views on the risky assets) to overcome many of the well-known problems with mean-variance analysis. He and Litterman [1999] provide a complete economic interpretation of the framework, extended by Satchell and Scowcroft [2000]. Idzorek [2004] introduces an alternative expression of confidence in investor’s views. Fabozzi, Focardi, and Kolm [2006] offer different examples and trading strategies suitable for a variety of active market participants. Several extensions of the model have also been proposed: Herold [2003] uses a benchmark instead of the market portfolio as the base from which investors tilt away according to their views; Martellini and Ziemann [2007] extend the model to include higher moments and, like Da Silva et al. [2009], employ the model in an active management framework; and Zhou [2009] extends the normal data generation process.

EXCHANGE TRADED VOLATILITY PRODUCTS

The first exchange traded volatility products were futures on the S&P 500 implied volatility index (VIX), first listed on the CBOE futures exchange (CFE) in March 2004. VIX is calculated and distributed on a real-time basis by the Chicago Board Options Exchange (CBOE). A detailed explanation of the calculation and characteristics of VIX is given by Whaley [2000, 2009]. Futures for implied volatility on the Dow Jones Industrial Average index and the Nasdaq-100 index were also launched and traded on CFE but failed to reach the success of their S&P 500 counterpart and were eventually delisted. In Europe, the volatility market is dominated by futures and options on the EUROSTOXX 50 implied volatility index, VSTOXX. However, in terms of directional exposure to volatility, VIX futures remain the focus for both investment and research. The success of VIX futures contracts led to the introduction of VIX options on the CBOE in February 2006. Trading volatility options allows trades that are both complex (e.g. by exposure to volatility-of-volatility rather than directional trades alone) and highly leveraged.

Early in 2009 the first, broker-traded volatility exchange traded notes (ETNs) were issued by Barclays Bank PLC: VXX tracks the performance of a Short-Term VIX futures index and VXZ tracks a Mid-Term VIX futures index.¹ The ETNs are issued with a ten-year tenor and returns on each note are that of the respective index minus an investor fee. The launch of the Barclays' volatility ETNs precipitated a considerable amount of bad press surrounding the negative roll yield when the VIX futures term structure is in contango. Since contango is experienced much of the time – it is only during crisis periods that the volatility term structure briefly exhibits some backwardation – these volatility ETNs make small losses due to daily rebalancing. In addition, ETNs are directly affected by the credit risk of the issuer. In crash periods, where investors expect high returns from such an investment, they are also subject to higher credit risk on behalf of the issuer. Notwithstanding these shortcomings, volatility ETNs were an instant success, with VXX turning over an average of 1.7 million shares per day during its first six months of trading. Barclays issued a third ETN in July 2009, labelled XXV, that offers the inverse performance of the VXX, thus giving Barclays the unique advantage of a perfect hedge for VXX. The popularity of volatility ETNs was confirmed with the introduction of VXX and VXZ options that are now traded on CBOE. At the time of writing the landscape for trading volatility is changing very rapidly.²

DATA

Equity exposure is characterized by a long position in SPY and the 1-month U.S. Treasury bill is used as the risk-free rate.³ Index volatility exposure is through the VIX futures that are listed on CFE. From Bloomberg we obtained nearly six years of daily data on the open, high, low, close, bid and ask prices, as well as the daily volume and open interest levels for the full set of VIX futures contracts that traded between 26 March 2004 (the inception of the VIX futures contract) and 31 December 2010.

Three VIX futures series based on different rollover methodologies are constructed as follows: The 1-month and 3-month series commence with the prompt futures contract on the first day of the sample. On the rollover date the algorithm chooses the next available contract (in the 1-month series) or the next available contract on the March quarterly cycle (the 3-month series). The third series, termed the Longest Traded (LT) series, always rolls over into the longest maturity contract that is actively traded.⁴ Rollovers are performed either 5 business days prior to a contract's maturity or upon maturity. Rollover usually induces a negative return because the VIX futures market is typically in contango; backwardation is experienced only during a period of unusually high volatility.

Exhibit 1 depicts the daily evolution of the values of SPY and the three VIX futures series rolling over 5 days prior to expiry. One can clearly observe the strong negative correlation between equity and volatility futures, particularly during the banking crisis when the upward trend of VIX futures prices could have preserved the positive returns on a long equity exposure.

Following Bessembinder and Chan [1992] returns on the three VIX series are defined by observations on the same contract, and the return on rollover day is based on the open price of the new contract rather than the previous day closing price. Since futures are instruments that require zero investment to enter the transaction special care should be taken on the calculation of VIX futures returns. Following Bodie and Rosansky [1980], an investor who takes a long position in a futures contract agrees to pay its current price (P_t) at the end of the investment period or equivalently the present value of the current price $W_t = P_t/(1 + R_{ft})$ now, where R_{ft} denotes the current risk-free rate of return prevailing for that period. If we denote by P_{t+1} the futures price at the end of the investment period and $R_t = (P_{t+1} - P_t)/P_t$ the percentage change in the futures price, then the return earned on W_t is:

$$\tilde{R}_t = \frac{P_{t+1}(1 + R_{ft})}{P_t} - 1 = (1 + R_t)(1 + R_{ft}) - 1 \approx R_t + R_{ft} \quad (1)$$

Thus, the percentage change in a futures price actually represents the excess return above the current risk-free rate earned on a futures position. The same result is used in a wide range of studies, including Fortenbery and Hauser [1990], Erb and Harvey [2006], Gorton, Hayashi, and Rouwenhorst [2008] and Daskalaki and Skiadopoulos [2010]. The above approximation becomes exact under continuously compounded (log) returns.

Exhibit 2 reports descriptive statistics for discrete and log returns: excess returns on SPY and ordinary returns on the three VIX futures series, for two sub-sample periods of roughly equal size that capture different regimes in the market: from 26 March 2004 – 31 May 2007 (a tranquil period) and 1 June 2007 – 30 June 2010 (a more volatile period). The sample is completed with a third sub-period from 1 July 2010 – 31 December 2010 which has features of the same nature (but greater in magnitude) as period 1.⁵ Comparing the results for sub-periods 1 and 2, SPY had a negative average excess return in period 2 whilst returns on VIX futures were high and positive on average, both series have higher volatility in period 2 and the negative correlation between SPY returns and VIX returns is also more pronounced than in period 1.⁶ This strengthens the view that a mix of equity and volatility during period 2 could have been beneficial to the long equity investor. Also note that log returns typically underestimate discrete returns.⁷ Despite the theoretical tractability of log returns, returns on an investment are discrete and so these are used in the subsequent analysis, which is the usual practice in asset allocation problems.

EMPIRICAL CHARACTERISTICS OF VIX FUTURES

The first step towards an efficient investment is to analyse liquidity, transaction costs and equilibrium returns. Whilst volatility may offer an excellent candidate for downside protection, a wide range of empirical studies demonstrate that the volatility risk premium is strongly negative.⁸ In order to explore these issues we analyse the trading volume of the VIX futures market, the bid-ask spread and the cost of carrying a long position in a VIX futures contract. We seek an optimal rollover strategy that will take into account these costs in deciding the frequency of rebalancing.

High transaction costs and low trading volumes used to be a significant factor that eroded gains from volatility trades. However, liquidity has improved significantly since the banking crisis. To see this, Exhibit 3 depicts the daily bid-ask spread and trading volume for the 1-month rollover series, exponentially smoothed to provide greater clarity of trends.⁹ Although still much higher than spreads on S&P 500 futures, the bid-ask spread on VIX futures has clearly diminished, and trading volume has significantly in-

creased during the last few years. The CFE extended trading hours on VIX futures in early December 2010, following an enormous increase in trading volume during the recent economic crisis in the Eurozone (last quarter of 2010). This, accompanied by a marked decrease in transaction costs, implies that many fund managers might now be trusting volatility products to provide downside protection against long equity exposures.

A prevalent characteristic of the volatility market is its negative carry (here ‘carry’ is the discrete return on a buy-and-hold strategy in a VIX futures contract). Exhibit 4 shows the carry that is realised on the 1-month series at each rollover, comparing the carry when the contract is rolled over upon maturity ($t = 0$) with a rollover date 5 business days before expiry ($t = 5$). All carry results are translated into a monthly equivalent (22 trading days) for better comparison. It is evident that most contracts had negative carry until the banking crisis, when VIX futures contracts produced very large positive returns. Moreover the vast majority of contracts have a lower and less variable carry when rolled over 5 days prior to maturity, rather than waiting until maturity. This is due to the familiar maturity effect, which is exacerbated by the settlement process for VIX futures. The underlying VIX index is based on the average of bid and ask prices of options entering the calculation formula but VIX futures are settled on the special opening quotation (SOQ) price. The SOQ is extracted using *actual* traded prices of SPX options during the market open at settlement day. Consequently, the difference between the VIX futures settlement price and VIX open deviates from zero. This convergence problem leads to increased arbitrage trading activity over the last few days prior to expiry, causing increased volatility in VIX futures prices as they approach expiry. The finding that rollover is not optimal upon maturity is consistent with other evidence on optimal rollover in futures markets, see e.g. Ma, Mercer, and Walker [1992]. From henceforth we shall rollover 5 days prior to expiry in all our empirical results.

Exhibit 5 presents the average spread, volume and total carry from all three VIX futures rollover strategies over the different sub-samples. The total carry represents the percentage growth of an investment in the three VIX futures series, initiated at the beginning of each sub-sample and held until it ends. Clearly, all three series experience a significant reduction in the bid-ask spread over time. The Longest Traded (LT) series carries the highest transaction costs and the most uncertain spread. However, whilst one gains from the 1 month rollovers in terms of higher liquidity and smaller spreads, the more frequent rollovers actually result in higher total costs than the other two strategies. In fact, the less frequent rollover in the LT series results in a higher overall total return.

JUSTIFICATION FOR VOLATILITY DIVERSIFICATION

This section first presents an ex-post empirical analysis of the diversification benefits from adding volatility to a positive exposure on US equities. More precisely, we ask whether it has ever been optimal to add a long VIX futures position to a long position on the SPY within the Markowitz [1952] framework. Then we ask: given that an investor is long SPY, how large does the expected return on VIX futures need to be in order to justify adding a long position on VIX to the SPY portfolio? We address this question using the rollover days for the three VIX futures series as the rebalancing points for an ex-ante portfolio construction. The problem is first approached in the mean-variance framework and then we employ the Black and Litterman [1992] model which allows investors to account for equilibrium returns on risky assets and incorporate personal views in the allocation problem. In each case historical inputs are based on two different in-sample periods, namely 1 month and 3 months.

Ex-post Mean-Variance Optimality

In this case the allocation of funds between two risky assets (SPY and VIX futures) and a risk-free asset (discount bond) may be considered in two stages: (i) find all mean-variance efficient combinations of SPY and VIX futures; and (ii) find the optimal mix of one of these portfolios with the risk-free asset. The portfolio chosen from stage (i) analysis is the tangency portfolio that when connected with the risk-free asset yields a linear efficient frontier with slope equal to the maximized Sharpe ratio (SR); the optimal choice along this frontier in stage (ii) is the portfolio that maximizes investor's expected utility. To answer the question posed above we need not consider specific risk preferences. We are only concerned with the convex frontier in {expected return, standard deviation} space in stage (i). In particular, we seek the times and conditions under which a mean-variance optimal allocation that is positive to SPY would also be positive to VIX futures. Thus, our problem may be stated mathematically in terms of the constrained optimization problem:

$$\max_w SR = \frac{w' \mu}{\sqrt{w' \Sigma w}}, \quad w_s + w_v = 1, \quad w_s > 0, w_v \geq 0, \quad (2)$$

where $w = (w_s, w_v)'$, μ is the column vector of expected excess return on SPY and expected return on VIX futures and Σ is the covariance matrix of SPY excess returns and VIX futures returns.

Our results begin by considering an ex-post optimal investment in the June 2010 VIX futures contract. Running from 22 October 2009 to 16 June 2010, it commences during a relatively tranquil period but matures just after the recent Greek crisis. For the period 22 October 2009 – 31 March 2010 the solution to (2) is achieved with $w_v = 0$. It only becomes optimal to add a positive VIX futures position on 1 April 2010, and hold it until 9 June 2010, a few days before the contract expires. This portfolio allocates 70.45% to SPY and 29.55% to the June 2010 VIX futures and would have produced an annualised SR of 3.61. By contrast, during this period a long position on the SPY alone had a negative SR of -1.89.

Empirical results extended over a longer period are based on the three VIX futures series which are defined by different rollover strategies, as explained in the previous section. It is again convenient to separate the total sample into sub-samples having different characteristics. In periods 1 (26 March 2004 - 31 May 2007) and 3 (1 July 2010 – 31 December 2010) the solution to (2) is achieved with $w_v = 0$. Thus it is clear that, despite the negative correlation between equity and volatility, a strong negative volatility risk premium is apparent when equity markets are relatively stable and trending, and it is never optimal to diversify into volatility. However, in period 2 (1 June 2007 - 30 June 2010) the SR of the portfolio is substantially greater when the long VIX futures positions summarized in Exhibit 6 are taken. During this period the ex-post optimal VIX exposure was always greater than 50% of the allocation to risky assets. It varied from 51.91% using the 3-month rollover, to 58.02% using the Longest Traded (LT) strategy and 63.45% using the 1-month rollover.

The highest SR is achieved by a VIX investment based on the LT rollover strategy, i.e. rolling over into the longest VIX futures maturity possible, provided it is reasonably liquid. If, on 1 June 2007, an investor had allocated 58.02% of the funds earmarked for SPY into VIX futures instead, and he had employed the LT rollover strategy, he would have achieved an average annualized SR of 0.8290 over the whole three years. This is greater than the SR obtained using the other two rollover strategies and much greater than a SR of -0.2996, which is what he would have obtained had he invested in SPY alone.

Exhibit 7 depicts the evolution over three years of a theoretical investment, on 1 June 2007, of (a) \$100 in the optimal combination of SPY and LT strategy, and (b) \$100 in SPY alone. Clearly, an equity-volatility diversified portfolio would have performed notably better. The difference between the performance of the two portfolios increases significantly during the credit and banking crisis and only reduces during the first half of

2010. However, the diversified portfolio again performs particularly well during the recent Greek crisis. During the three-year period, the SPY-VIX portfolio produced an annualized excess return of 17.31% return with 20.89% volatility compared to an annualized negative excess return of -9.03% with 30.15% volatility for the SPY alone.

The maximization of SR is equivalent to the mean-variance criterion when returns are normally distributed and investors have an exponential utility function. In this case, maximizing expected utility is equivalent to maximizing the certainty equivalent (CE) return, which is given by $w'\mu - \frac{1}{2}\gamma\sigma^2$, where $\sigma^2 = w'\Sigma w$ and γ denotes the investor's level of risk aversion. A proof is provided in Freund [1956]. The resulting maximized value of the CE return is a function of the SR. However, there is strong empirical evidence that returns on financial assets are not normal. We therefore extended our investigation to include the impact of non-zero skewness, τ and excess kurtosis, κ . In this case the optimal allocation will be that which maximizes the generalised Sharpe ratio (GSR) defined as $(-2\ln(-E(U)^*))^{1/2}$, where $E(U)^*$ denotes the maximum expected utility – see Hodges [1998]. For an investor with exponential utility we may use the following approximation:

$$GSR \approx (SR^2 + \frac{\tau}{3}SR^3 - \frac{\kappa}{12}SR^4)^{1/2}, \quad (3)$$

as shown in Alexander [2008a]. For brevity, we do not report here the portfolios obtained when maximizing (3), save to say that they had very similar characteristics to those reported above. The important point to note is that, even without an inappropriate normality assumption, it is ex-post optimal to add volatility to a long equity exposure during excessively volatile periods.

Ex-ante Mean-Variance Analysis

When an investor allocates between SPY, VIX futures and a discount bond according to the mean-variance criterion, maximizing the CE return yields the unconstrained solution $w = \gamma^{-1}\Sigma^{-1}\mu$, where $w = (w_s, w_v)'$. In this sub-section the problem is not considered in two stages, as it was in the ex-post analysis, so w_s and w_v are not constrained to sum to 1, the allocation being completed with the residual invested in the risk-free asset. Let σ_s^2 , σ_v^2 and σ_{sv} denote the SPY excess returns variance, VIX futures returns variance and their covariance, respectively. Then the unconstrained solution may be written:

$$w_s = \frac{1}{\gamma|\Sigma|}(\sigma_v^2\mu_s - \sigma_{sv}\mu_v) \quad , \quad w_v = \frac{1}{\gamma|\Sigma|}(\sigma_s^2\mu_v - \sigma_{sv}\mu_s),$$

where $|\Sigma| = \sigma_s^2 \sigma_v^2 - \sigma_{sv}^2$ is the determinant of Σ . Since $\gamma > 0$, $|\Sigma| > 0$ and $\sigma_{sv} < 0$, requiring both $w_s > 0$ and $w_v > 0$ simultaneously results in the conditions:

$$\mu_v > \frac{\mu_s \sigma_v^2}{\sigma_{sv}} \quad \text{and} \quad \mu_v > \frac{\mu_s \sigma_{sv}}{\sigma_s^2}.$$

Equivalently, $w_s > 0$ and $w_v > 0$ is guaranteed when

$$\mu_v > \max \left(\frac{\mu_s \sigma_v^2}{\sigma_{sv}}, \frac{\mu_s \sigma_{sv}}{\sigma_s^2} \right). \quad (4)$$

Note that the inequality (4) does not depend on the level of risk aversion.

The above may be used to investigate how large the expected return on VIX futures needs to be to justify adding a long position on VIX to the SPY portfolio on each rebalancing date. Now μ_s , σ_s^2 , σ_v^2 and σ_{sv} are set equal to their historical (in-sample) values and the minimum expected return on VIX futures is derived using (4). Exhibit 8 presents the results based on the 1-month rollover futures investment strategy with a 1-month in-sample period. This indicates that a long equity investor, who allocates in a mean-variance optimal fashion with an historical covariance matrix based only on recent market data, could very often justify diversification into VIX futures. In fact, based on our assumptions it could be optimal to diversify into volatility even when it has a very large and negative expected return. For instance, in April 2009 an expected return of -27% over the next month would be sufficient to justify diversification into VIX futures.¹⁰

However, this conclusion should be regarded with caution because it relies heavily on the use of an in-sample mean for μ_s and a covariance matrix estimate based on only recent historical data. The effect of errors in these estimates has been thoroughly discussed in the literature – see Best and Grauer [1991] and Chopra and Ziemba [1993] – and it is errors in the mean returns that have the predominant effect on results. In any case, a long-term investment should take account of equilibrium expected returns, as in the Black and Litterman [1992] model, and such investors should not base their decisions on short-term expectations alone.

Ex-ante Black-Litterman Analysis

We now answer the same question as above using the classical Black-Litterman model as in He and Litterman [1999]. First suppose that asset's returns follow a normal distribution with unknown expected returns vector μ and covariance matrix Σ . In equilibrium,

all investors hold the same view for expected returns, expressed as the equilibrium risk premium vector Π . Then the Bayesian prior for expected returns is a normal distribution with mean Π and covariance matrix $\tau\Sigma$, where τ is a constant indicating the uncertainty in the prior. Additional to the prior beliefs an investor holds current views on the asset's returns. These are expressed via a matrix P such that $P\mu$ follows a normal distribution with mean vector Q and diagonal covariance matrix Ω ; Q contains investor's current expected returns and Ω describes the confidence in each view. Blending equilibrium with current views yields a posterior normal distribution for expected returns with mean

$$\mu_{BL} = [(\tau\Sigma)^{-1} + P'\Omega^{-1}P]^{-1} [(\tau\Sigma)^{-1}\Pi + P'\Omega^{-1}Q] \quad (5)$$

and covariance matrix $M = [(\tau\Sigma)^{-1} + P'\Omega^{-1}P]^{-1}$. Then the asset's returns follow a normal distribution with mean (5) and covariance matrix $\bar{\Sigma} = \Sigma + M$. Finally, He and Litterman [1999] use the mean-variance optimizer to obtain the solution for the unconstrained optimal portfolio weights:

$$w^* = \frac{1}{1 + \tau}(w_{eq} + P'\Lambda),$$

where $w_{eq} = \gamma^{-1}\Sigma^{-1}\Pi$ are the equilibrium portfolio weights and vector Λ is given by:

$$\Lambda = \gamma^{-1}\tau\Omega^{-1}Q - A^{-1}P\frac{\Sigma}{1 + \tau}w_{eq} - A^{-1}P\frac{\Sigma}{1 + \tau}P'\gamma^{-1}\tau\Omega^{-1}Q, \quad (6)$$

with $A = \tau^{-1}\Omega + P\frac{\Sigma}{1 + \tau}P'$.

We use this framework to solve for the current expected return or view on volatility that justifies the addition of VIX futures to the SPY portfolio. Current views are expressed directly on the SPY and VIX returns, i.e. P equals the identity matrix. The elements of Q are denoted q_s and q_v , for SPY and VIX respectively, and we set q_s equal to the historical (in-sample) mean. Following the suggestion of He and Litterman [1999] we express confidence in our current views using $\Omega = \text{diag}(P\tau\Sigma P')$ and Σ is the historical covariance matrix with elements σ_s^2 , σ_v^2 and σ_{sv} . Substituting in (6) leads, after some straightforward algebra, to the following expression for $\Lambda = (\lambda_s, \lambda_v)'$:

$$\lambda_s = \frac{q_s\sigma_v^2(\tau + 2)(\tau + 1) - q_v(\tau + 1)\sigma_{sv} - \mu_s(\tau + 2)\sigma_v^2 + \mu_v\sigma_{sv}}{\gamma [(\tau + 2)^2\sigma_s^2\sigma_v^2 - \sigma_{sv}^2]},$$

$$\lambda_v = \frac{q_v \sigma_s^2 (\tau + 2) (\tau + 1) - q_s (\tau + 1) \sigma_{sv} - \mu_v (\tau + 2) \sigma_s^2 + \mu_s \sigma_{sv}}{\gamma [(\tau + 2)^2 \sigma_s^2 \sigma_v^2 - \sigma_{sv}^2]}.$$

The solution can now be written:

$$w_s^{BL} = \frac{1}{1 + \tau} \left(\frac{1}{\gamma |\Sigma|} (\sigma_v^2 \mu_s - \sigma_{sv} \mu_v) + \lambda_s \right), \quad w_v^{BL} = \frac{1}{1 + \tau} \left(\frac{1}{\gamma |\Sigma|} (\sigma_s^2 \mu_v - \sigma_{sv} \mu_s) + \lambda_v \right).$$

Since $\gamma > 0$, $|\Sigma| > 0$ and $\sigma_{sv} < 0$, requiring both $w_s > 0$ and $w_v > 0$ yields:

$$q_v > \frac{-\mu_s^{eq} \sigma_v^2 [(\tau + 2) \sigma_s^2 \sigma_v^2 + \sigma_{sv}^2] + \mu_v^{eq} \sigma_s^2 \sigma_v^2 \sigma_{sv} (\tau + 3) - q_s \sigma_v^2 (\tau + 2) [\sigma_s^2 \sigma_v^2 - \sigma_{sv}^2]}{-\sigma_{sv} [\sigma_s^2 \sigma_v^2 - \sigma_{sv}^2]} = a_1$$

$$q_v > \frac{\mu_s^{eq} \sigma_s^2 \sigma_v^2 \sigma_{sv} (\tau + 3) - \mu_v \sigma_s^2 [(\tau + 2) \sigma_s^2 \sigma_v^2 + \sigma_{sv}^2] + q_s \sigma_{sv} [\sigma_s^2 \sigma_v^2 - \sigma_{sv}^2]}{\sigma_s^2 (\tau + 2) [\sigma_s^2 \sigma_v^2 - \sigma_{sv}^2]} = a_2$$

Equivalently, $w_s > 0$ and $w_v > 0$ is guaranteed when $q_v > \max(a_1, a_2)$. Again the inequality does not depend on the individual investor's level of risk aversion.

The equilibrium return on SPY is set at $\mu_s^{eq} = 4\%$ per annum and that on VIX futures is set at $\mu_v^{eq} = -40\%$. This choice is consistent with seminal work on the equity risk premium by Fama and French [2002] and the volatility risk premium by Carr and Wu [2009] among others. We set $\tau = 0.05$, this value also being consistent with standard practice. Now taking account of the equilibrium returns on the two assets, and in particular the large negative volatility risk premium, we use the above inequality to re-calculate the minimum short-term expected return on VIX futures that justifies a positive investment in both SPY and VIX futures.

Recall that under the mean-variance approach long equity investors may find diversification into volatility optimal even when their expected returns on it are negative. This is no longer true in the Black-Litterman approach. Due to the large negative volatility risk premium, views on VIX futures expected returns always need to be positive (and are often quite large) to justify volatility diversification for a long equity investor.

We further examine these Black-Litterman views on volatility, asking: is it reasonable for a long equity investor to hold such views? To answer this requires a model of expectation formation and we consider two simple extremes, static expectations and perfect foresight. Under static expectations the investor forms views based entirely on recent historical data, i.e. the expected return is equal to the in-sample mean. Under perfect foresight the investor's expected return is correct, i.e. set equal to the out-of-sample realised return, which can only be computed ex-post.

Exhibit 9 presents the correlations between the Black-Litterman views on volatility that justify its use for diversification, with (a) the historical returns on VIX during the recent past and (b) the realised returns until the next rebalancing point.¹¹ The significant positive correlations with historical returns implies that investors with static expectations could very often justify diversification into VIX futures in the Black-Litterman framework. This adds weight to our previous results based on the mean-variance framework.

However, investors with perfect foresight would rarely consider volatility diversification to be optimal in the Black-Litterman model. We see this because the correlations between Black-Litterman views and realised returns in Exhibit 9 are small and insignificantly different from zero. This implies that, whilst it is easy to justify volatility diversification in an optimization framework, volatility diversified portfolios have an out-of-sample performance that is typically worse than holding equity alone.

To illustrate this we computed the SR for two portfolios in the Black-Litterman setting described above: one that allocates only between SPY and the discount bond; and one that considers diversification into VIX futures, based on static expectations with a 3-month in-sample period. The LT rollover strategy is employed to minimize the impact of negative roll yield. We assumed a typical value of risk aversion ($\gamma = 4$) and examined the relative out-of-sample performance of these portfolios over approximately 6 years, between November 2004 and December 2010.

A constant holding on SPY had a SR that was very low (0.10) over this period. The Black-Litterman optimally-rebalanced portfolio did much better with a SR of 0.21, provided the allocation was between SPY and discount bond alone. However, the diversified portfolio which included the possibility of investing in VIX futures had a SR of only 0.15. Further results are not reported for brevity. It suffices to say that their qualities are robust to our choice of parameters and rollover strategy above. They clearly show that adding VIX futures to optimal equity portfolios will deteriorate long-term performance.

CONCLUSION

Volatility trading is becoming very popular. However, long equity institutional investors such as mutual and pension funds should be wary of regarding volatility as a useful diversification tool. Only experienced investors should consider trading it, using volatility products such as futures, ETNs and variance swaps for directional exposure to spot and/or forward volatility, arbitrage trading and vega hedging. Whilst market makers, prop traders, hedge fund analysts, structurers in large banks and other sophisticated

market players should have the necessary expertise to speculate and hedge with volatility, it is unlikely that institutional portfolio managers have the sophistication needed to enter this market, especially when the pensions of ordinary people are at stake.

Justifying the sale of VIX futures to institutional investors is easy in both the mean-variance and Black-Litterman frameworks (for instance, using static volatility expectations based on historical VIX futures returns). Equity portfolio managers may even justify buying a VIX futures contract when it is expected to have a large negative return, provided they ignore its equilibrium return. Moreover, an ex-post analysis could demonstrate that buying VIX futures would considerably enhance equity returns, provided the sample was between June 2007 and June 2010.

Unfortunately, this story is rather misleading. Perfect foresight would seldom justify the purchase of VIX futures as a long equity diversification tool. Most of the time volatility's negative carry and roll yield heavily erodes equity performance, and the only time volatility diversification is optimal is at the onset of a stock market crisis. The problem is that such crises are extremely difficult to predict and relatively short-lived. In other words, equity volatility is characterised by unexpected jumps followed by very rapid mean-reversion, so expectations based on recent volatility behaviour are unlikely to be realised. In short, by the time we are aware of a crisis it is usually too late to diversify into volatility.

ENDNOTES

¹The short-term index trades the two near-term VIX futures contracts and rolls continuously from the first to the second month in order to obtain a constant one-month maturity exposure. Similarly, the mid-term index uses the fourth, fifth, sixth, and seventh month VIX futures contracts and captures a constant five-month maturity exposure.

²For instance, between late 2010 and early 2011 several similar ETNs were launched by rival companies. Notably, Credit Suisse launched six new volatility ETNs on behalf of VelocityShares, some that compete directly with the Barclay's offerings and others that offer $2 \times$ leverage. Citi issued its own version of a volatility ETN (CVOL), which is linked to an index that offers a long exposure on the third and fourth month VIX futures contracts coupled with a short S&P 500 position. Barclays responded with the introduction of VZZ which offers a leveraged return on the S&P 500 VIX Mid-Term futures index, but also moved on to diversified ETNs such as VQT, that tracks the S&P 500 Dynamic VEQTOR index which is essentially a portfolio that allocates dynamically between long equity (S&P 500 index), long volatility (S&P 500 VIX Short-Term futures index) and cash (overnight LIBOR).

³SPY is the code for Standard and Poor's Depository Receipts, an exchange traded

fund (ETF) which replicates the performance of the S&P 500 index and trades in the NYSE Arca Exchange. The 1-month TBill is used because our most frequently traded rollover strategy rebalances monthly. Daily closing prices of the SPY were obtained from Datastream and the 1 month U.S. Treasury bill rates were downloaded from the U.S. Federal Reserve website.

⁴We assume any futures traded must have a 10-day average volume of at least 50 contracts.

⁵In period 3, and under discrete returns (log returns), SPY exhibited an annualized average excess return of 39.62% (37.46%) and the three VIX futures exhibited returns of -234.19% (-248.99%), -144.87% (-155.29%), and -104.50% (-110.01%), respectively. Annual standard deviations are 15.03% (15.06%) for SPY and 52.15% (52.50%), 44.60% (44.90%), and 32.48% (32.64%) for the three VIX futures series. Given these large and negative average VIX returns, the case against the diversification power of VIX futures during periods with low volatility is supported.

⁶The correlations on discrete returns range between -0.68 (for the Longest Traded VIX series in period 1) and -0.78 (for the 1-month VIX series in period 2). All correlations are of greater magnitude in period 2, and are greater for the VIX series with more frequent rollover.

⁷In fact, if $x = P_{t+1}/P_t$ then $\ln(x) \leq x - 1$, for $x > 0$ and $x - 1 \leq 1$. We must have some $P_t \leq 0$ or some $P_{t+1} > 2P_t$ for the respective conditions to be violated. This is usually not the case in our database, nor in most empirical time series of financial asset returns.

⁸Such evidence can be found, among others, in Bakshi and Kapadia [2003a,b] and Carr and Wu [2009].

⁹The bid-ask spread is defined as the absolute difference of the ask and bid sides of the market as percentage of the mid price while trading volume figures represent the number of contracts traded in the market on any particular day.

¹⁰Qualitatively similar results are obtained based on a 3-month in-sample period, and for the other two VIX futures series based on both in-sample periods, so these are not reported for brevity.

¹¹The historical and realised returns are the average daily return over the relevant period, converted into a monthly equivalent for the 1-month series, a 3-month equivalent for the 3-month series and a 6-month equivalent for LT. We also derived results using the total return over the respective period, but the results were practically identical and are not reported for brevity.

REFERENCES

- Alexander, C. *Market Risk Analysis Volume I: Quantitative Methods in Finance*. Chichester: Wiley, 2008a.
- . *Market Risk Analysis Volume II: Practical Financial Econometrics*. Chichester: Wiley, 2008b.
- Bakshi, G., and N. Kapadia. “Delta-Hedged Gains and the Negative Market Volatility Risk Premium.” *Review of Financial Studies*, 16, 2 (2003a), pp. 527–566.
- . “Volatility Risk Premiums Embedded in Individual Equity Options.” *The Journal of Derivatives*, 11, 1 (2003b), pp. 45–54.
- Bessembinder, H., and K. Chan. “Time-varying risk premia and forecastable returns in futures markets.” *Journal of Financial Economics*, 32, 2 (1992), pp. 169–193.
- Best, M., and R. Grauer. “On the sensitivity of mean-variance-efficient portfolios to changes in asset means: some analytical and computational results.” *Review of Financial Studies*, 4, 2 (1991), pp. 315–342.
- Black, F., and R. Litterman. “Global portfolio optimization.” *Financial Analysts Journal*, 48, 5 (1992), pp. 28–43.
- Bodie, Z., and V. Rosansky. “Risk and return in commodity futures.” *Financial Analysts Journal*, 36, 3 (1980), pp. 27–39.
- Carr, P., and L. Wu. “Variance risk premiums.” *Review of Financial Studies*, 22, 3 (2009), pp. 1311–1341.
- Chen, H., S. Chung, and K. Ho. “The diversification effects of volatility-related assets.” *Journal of Banking and Finance*, 2010.
- Chopra, V., and W. Ziemba. “The effect of errors in means, variances, and covariances on optimal portfolio choice.” *The Journal of Portfolio Management*, 19 (1993), pp. 6–11.
- Da Silva, A., W. Lee, and B. Pornrojngkool. “The Black–Litterman model for active portfolio management.” *The Journal of Portfolio Management*, 35, 2 (2009), pp. 61–70.
- Daigler, R., and L. Rossi. “A portfolio of stocks and volatility.” *The Journal of Investing*, 15, 2 (2006), pp. 99–106.

- Dash, S., and M. Moran. “VIX as a companion for hedge fund portfolios.” *The Journal of Alternative Investments*, 8, 3 (2005), pp. 75–80.
- Daskalaki, C., and G. Skiadopoulos. “Should Investors Include Commodities in Their Portfolios After All? New Evidence.” Working Paper, 2010.
- Egloff, D., M. Leippold, and L. Wu. “Variance risk dynamics, variance risk premia, and optimal variance swap investments.” *Journal of Financial and Quantitative Analysis (forthcoming)*, 2010.
- Erb, C., and C. Harvey. “The strategic and tactical value of commodity futures.” *Financial Analysts Journal*, 62, 2 (2006), pp. 69–97.
- Fabozzi, F., S. Focardi, and P. Kolm. “Incorporating trading strategies in the Black-Litterman framework.” *The Journal of Trading*, 1, 2 (2006), pp. 28–37.
- Fama, E., and K. French. “The equity risk premium.” *Journal of Finance*, 57, 2 (2002), pp. 637–659.
- Fortenbery, T., and R. Hauser. “Investment potential of agricultural futures contracts.” *American Journal of Agricultural Economics*, 72, 3 (1990), pp. 721–726.
- Freund, R. “The introduction of risk into a programming model.” *Econometrica: Journal of the Econometric Society*, 24, 3 (1956), pp. 253–263.
- Gorton, G., F. Hayashi, and K. Rouwenhorst. “The fundamentals of commodity futures returns.” Yale ICF Working Paper, 2008.
- Hafner, R., and M. Wallmeier. “Optimal Investments in volatility.” *Financial Markets and Portfolio Management*, 22, 2 (2008), pp. 147–167.
- He, G., and R. Litterman. “The intuition Behind the Black-Litterman Model Portfolios.” Goldman Sachs Investment Management Series, 1999.
- Herold, U. “Portfolio construction with qualitative forecasts.” *The Journal of Portfolio Management*, 30, 1 (2003), pp. 61–72.
- Hodges, S. A generalization of the Sharpe ratio and its applications to valuation bounds and risk measures. Financial Options Research Centre, Warwick Business School, Univ. of Warwick, 1998.

- Idzorek, T. “A step-by-step guide to the Black-Litterman model.” Working Paper, 2004.
- Lee, B., and Y. Lin. “Using Volatility Instruments as Extreme Downside Hedges.” Working paper, 2010.
- Ma, C., J. Mercer, and M. Walker. “Rolling over futures contracts: A note.” *Journal of Futures Markets*, 12, 2 (1992), pp. 203–217.
- Markowitz, H. “Portfolio Selection.” *Journal of Finance*, 7, 1 (1952), pp. 77–91.
- Martellini, L., and V. Ziemann. “Extending Black-Litterman analysis beyond the mean-variance framework.” *The Journal of Portfolio Management*, 33, 4 (2007), pp. 33–44.
- Moran, M., and S. Dash. “VIX Futures and Options.” *The Journal of Trading*, 2, 3 (2007), pp. 96–105.
- Pézier, J., and A. White. “The relative merits of alternative investments in passive portfolios.” *The Journal of Alternative Investments*, 10, 4 (2008), pp. 37–49.
- Satchell, S., and A. Scowcroft. “A demystification of the BlackLitterman model: Managing quantitative and traditional portfolio construction.” *Journal of Asset Management*, 1, 2 (2000), pp. 138–150.
- Szado, E. “VIX Futures and Options: A Case Study of Portfolio Diversification During the 2008 Financial Crisis.” *The Journal of Alternative Investments*, 12 (2009), pp. 68–85.
- Whaley, R. “The investor fear gauge.” *The Journal of Portfolio Management*, 26, 3 (2000), pp. 12–17.
- . “Understanding VIX.” *The Journal of Portfolio Management*, 35 (2009), pp. 98–105.
- Zhou, G. “Beyond Black–Litterman: Letting the Data Speak.” *The Journal of Portfolio Management*, 36, 1 (2009), pp. 36–45.

EXHIBIT 1

SPY ETF and VIX Futures.

Time series of SPY (left scale) and the three VIX futures series (right scale)



EXHIBIT 2

Descriptive Statistics.

Summary statistics for the SPY (excess returns) and VIX futures (returns).

Discrete Returns								
Panel A: 26 March 2004 - 31 May 2007					Panel B: 1 June 2007 - 30 June 2010			
	SPY	VIX			SPY	VIX		
		1-month Rollover	3-month Rollover	Longest Traded		1-month Rollover	3-month Rollover	Longest Traded
Annualized Mean	8.31%	-85.30%	-47.59%	-55.85%	-9.03%	32.08%	37.24%	36.37%
Volatility	10.48%	42.52%	37.72%	25.80%	30.15%	71.53%	56.23%	48.36%
Skewness	-0.39	2.17	3.20	-0.53	0.41	0.90	0.71	0.82
Excess Kurtosis	1.74	22.35	35.13	5.32	8.67	2.04	1.81	3.99
Sharpe Ratio	0.79	-2.01	-1.26	-2.16	-0.30	0.45	0.66	0.75
Log Returns								
Panel A: 26 March 2004 - 31 May 2007					Panel B: 1 June 2007 - 30 June 2010			
	SPY	VIX			SPY	VIX		
		1-month Rollover	3-month Rollover	Longest Traded		1-month Rollover	3-month Rollover	Longest Traded
Annualized	7.76%	-94.25%	-54.47%	-59.27%	-13.56%	7.14%	21.69%	24.86%
Volatility	10.49%	41.71%	36.63%	26.00%	30.10%	70.27%	55.57%	47.82%
Skewness	-0.43	1.50	2.39	-0.71	0.12	0.69	0.54	0.58
Excess Kurtosis	1.85	15.46	24.72	6.09	7.77	1.68	1.55	3.53
Sharpe Ratio	0.74	-2.26	-1.49	-2.28	-0.45	0.10	0.39	0.52

EXHIBIT 3

Bid-Ask Spread and Volume on VIX (1-month Rollover)

The closing ask price minus the closing bid price as percentage of the closing mid price (left scale), and the number of contracts traded each day (right scale). Series are smoothed exponentially with a smoothing constant of 0.9.




EXHIBIT 4

Carry from Long Position on VIX Futures (1-month Rollover).

This is the discrete return from a buy-and-hold strategy on any VIX futures contract that is part of the 1-month rollover series. All results are transformed into a monthly equivalent (22 trading days). We show the carry from holding the contract until maturity ($t = 0$) and until 5 business days prior to expiry ($t = 5$).




EXHIBIT 5

Spread, Volume and Total Carry on VIX Futures.

In each sub-sample we report: the mean and standard deviation of the daily bid-ask spread; the mean of the daily number of contracts traded; and the total return on a buy-and-hold strategy, rebalanced according to our three VIX futures series.

	Panel A 26 March 2004 – 31 May 2007			Panel B 1 June 2007 – 30 June 2010			Panel C 1 July 2010 – 31 December 2010		
	1-month Rollover	3-month Rollover	Longest Traded	1-month Rollover	3-month Rollover	Longest Traded	1-month Rollover	3-month Rollover	Longest Traded
Average Spread	0.77%	0.73%	0.73%	0.33%	0.44%	0.58%	0.24%	0.25%	0.27%
Spread Stdev	0.41%	0.39%	0.38%	0.25%	0.34%	0.54%	0.09%	0.10%	0.13%
Average Volume	390	356	354	2804	1520	1046	9356	6011	3287
Total Carry	-95.08%	-82.46%	-84.96%	24.84%	96.20%	116.55%	-72.85%	-54.85%	-43.06%

EXHIBIT 6

Ex-Post Mean-Variance Analysis for Optimal SPY-VIX Portfolios.

Optimal weights on SPY and VIX futures for a long equity investor. Sample period is 1 April 2010 - 9 June 2010 for the June 2010 VIX futures contract, and 1 June 2007 - 30 June 2010 for the three VIX futures series.

Contract	w_s	w_v	SR with VIX	SR SPY only
June 2010	70.45%	29.55%	3.6091	-1.8872
1m Rollover	36.55%	63.45%	0.4562	-0.2996
3m Rollover	48.09%	51.91%	0.7435	
Longest Traded	41.98%	58.02%	0.8290	

EXHIBIT 7

Portfolio Growth, 1 June 2007 - 30 June 2010.




EXHIBIT 8

Expected Returns on VIX Futures to Justify Diversification into Volatility (1-month Rollover).

The minimum expected monthly return on VIX required by a long equity investor at every rebalancing point to guarantee a positive VIX investment when allocating according to the mean-variance criterion. The in-sample period used is one month.




EXHIBIT 9

Diversification Efficiency

Correlations between Black-Litterman expected return on VIX futures that justify diversification into volatility for a long equity investor, with (a) the historical return, and (b) the realised return. Correlation t statistics were computed as in Alexander [2008b], p109.

	1-month in-sample			3-month in-sample		
	1m Roll-Over	3m Roll-Over	Longest Traded	1m Roll-Over	3m Roll-Over	Longest Traded
Historical	0.74	0.49	0.87	0.50	0.48	0.73
t-statistic	9.7354	2.9364	6.6392	5.0487	2.8742	4.0135
Realised	0.16	-0.06	-0.24	0.13	-0.04	0.00
t-statistic	1.3999	-0.3188	-0.9312	1.1245	-0.2167	0.0067