

The VIX-VIX Futures Puzzle*

Ivan Oscar Asensio[†]

October 2013

Abstract

Ex ante forecasts of future changes in the level of the VIX, implied by the VIX term structure, overshoot ex post changes, especially over shorter tenors¹. Strategies designed to profit from this phenomenon have not been successful in eliminating the arbitrage opportunity. In fact, the magnitude of the forecast bias has increased as additional capital has flowed into the strategy, a result which does not provide support for the slow-moving capital explanation of arbitrage persistence. Instead, I present evidence that suggests that the VIX-VIX Futures Puzzle is propagated by significant inflows of capital from non-professional investors that have entered VIX futures markets via the proliferation of ETF (exchange-traded fund) offerings.

*The VIX, short for Chicago Board of Exchange (CBOE) Volatility Index, represents the market's estimate of future realized volatility in the S&P 500 index over a one-month period. I am grateful for invaluable comments and suggestions from Michael Hutchison, Joshua Aizenman, Thomas Wu, Carl Walsh, Chris Limnios, 2013 ASSA Conference participants from LSU and Guggenheim Partners, and seminar attendees at UCSC, HSBC, and the University of Victoria Gustavson School of Business. Research support from the Cota-Robles Foundation and the Sury Initiative for Global Finance and International Risk Management (SIGFIRM) is also acknowledged.

[†]University of Victoria, Peter B. Gustavson School of Business, 3800 Finnerty Drive (Ring Road), Victoria BC V8P 5C2, phone: 212-402-7369, email: iasensio@uvic.ca.

¹Application of the expectations hypothesis (EH) suggests the slope of the VIX term structure today, constructed from VIX futures prices, should be consistent with the direction and magnitude of future changes of the VIX index itself. This is the analogue of the issue investigated by Shiller (1979) and Froot (1989) for the term structure of interest rates, by Campa and Chang (1995) for currency option implied volatilities, and by Mixon (2007) for equity index options.

1. Introduction

The application and relevance of implied volatility measures to asset pricing, risk management, investment portfolio construction, corporate accounting disclosure and monetary policy development reinforce the importance of evaluating its forecast accuracy. Two fundamental questions have been addressed in the literature. The first investigates the extent implied volatilities accurately forecast the ex post realized volatility of underlying asset prices. In other words, are option premiums incurred justified by the subsequent payouts on those options? The second fundamental question addressed in the literature evaluates the evolution of implied volatility itself across time. The existence of a term structure implies that the market assigns different prices for different time horizons. Does implied volatility evolve according to term structure forecasts? Specifically, when the term structure is positively (negatively) sloped, does implied volatility at the short-end of the curve rise (fall) as much as predicted? This paper falls under the latter mandate.

Implied volatility in this context, however, will not be derived from the prices of individual options but rather will be based on volatility products whose prices are derived according to a model-free approach. Fundamentally, model-free implied volatility can be derived from a portfolio of individual options with strikes spanning a defined range of possible outcomes for the underlying asset. Breeden and Lizenberg (1978), Demeterfi et al. (1999), Britten-Jones and Neuberger (2000), Carr and Wu (2006), Dupire (2006), and Jiang and Tian (2007) have described theoretical relationships between underlying option prices and model-free implied volatility, under the implicit assumption that the individual options on the underlying assets are priced properly.

The VIX is the most widely followed and well-known volatility product, representing the market's estimate of future volatility in the S&P 500 index (SPX) over a one-month period. The VIX has been embraced as a risk management vehicle by investors, a barometer for risk aversion by financial markets participants, and an input to econometric specifications and robustness tests by academic researchers. A higher reading on the VIX corresponds to greater aversion to market risk, while lower readings are associated with rising risk appetite. In fact, over the last few years the index has acquired global acceptance as the ultimate barometer for investor sentiment². The VIX is an index whose calculation is based on a set of option prices on the SPX across a wide variety of strikes, thus a negative co-movement between the VIX and the SPX exists by construction³. Figure 1 tracks historical price data on both indices, from January 2006 to October 2012. The strong association between both price series is supported by a correlation coefficient, based on daily

²The VIX index is commonly referred to as the market's 'fear gauge'.

³The co-movement is imperfect for reasons highlighted in section 3. Calculation of the VIX is outlined in this section.

price changes, of -76% for the entire period. In addition, the average three-month rolling correlation is -83%.

- Insert Figure 1 here -

Because the VIX is not an investable asset and because replication is costly, investment exposure to the index is achieved principally by trading VIX futures. The information content of futures markets is thus an important component for empirical work involving the VIX. The focus of this paper, then, is to evaluate of the expectations hypothesis (EH) on the VIX term structure⁴ based on the risk-neutral formulation proposed by Campa and Chang (1995). There are two main objectives. The first is to characterize and quantify the forecast bias, if any, implied by the VIX term structure. Does implied volatility evolve according to term structure forecasts? There are four key findings along these lines. First, VIX futures are consistently overpriced relative to the subsequent moves in the underlying VIX index. Second, the forecast bias increases substantially with shorter tenors. Third, the forecast bias is smaller comparatively during periods of extreme volatility. Four, deviations from expectations hypothesis are greatest during the years following the US credit crisis of 2008-9. These findings describe the VIX-VIX Futures Puzzle.

The second objective, once the forecast bias has been documented, is to shed light on why the forecast bias exists and why it persists. I first identify a number of factors that might influence the size of the forecast bias over time: futures open interest, hedge fund capital flows, performance of safe haven assets, the costs of insuring against tail risks, and the amount of credit risk in the financial system. I then construct an arbitrage strategy which aims to profit from the forecast bias reported, and evaluate the extent to which arbitrage profits are impacted by the set of aforementioned factors.

This paper is organized as follows. Section 2 provides a literature review. Section 3 defines the VIX and its term structure. Section 4 discusses the framework and theory that is used for evaluating EH. Section 5 describes the data. Section 6 contains the results of the various tests of expectations hypothesis. Section 7 discusses the various factors that drive the forecast bias. Section 8 addresses arbitrage existence and persistence. Section 9 concludes.

⁴Note, an innovation proposed in this work involves construction of the VIX term structure using actual VIX futures market prices.

2. Literature review

2.1 Implied volatility forecast bias

With regard to the forecast bias of implied volatility, there are two questions addressed in the literature.

2.1.A Implied volatility versus ex post realized volatility

The first strand of literature investigates whether option premiums are justified by the subsequent payouts on those options. Along these lines, Poon and Granger (2003) offer an thorough survey of volatility forecasting. Based on 93 peer-reviewed research articles, the authors establish that option implied volatility contains a significant amount of information about future realized volatility, often providing more accurate forecasts than model-based efforts derived from time series models⁵. Evidence is furnished for options on equities, broad-based price indices, interest rates, currencies, and commodities. Comparing across asset classes, it is reasonable to expect different levels of forecast accuracy for options written on different assets. This is primarily due to differences in trading frictions. Because market-makers rely on executing transactions in spot, forward, and futures markets as a way to manage the risk in their options inventory portfolios, assets which trade in the presence of lower frictions (and thus greater liquidity, narrower bid-ask spreads, etc.) would be expected to have options markets which exhibit lower levels of forecast bias. Along these lines, the literature has focused on testing implied volatility for at-the-money (ATM) strikes or utilizing weighting schemes across different strikes which overweight the ATM strikes. Liquidity is superior.

Buraschi and Jackwerth (2001), Coval and Shumway (2001), Bakshi and Kapadia (2002), and Pan (2002) all report that ex ante estimates derived from implied volatility overshoot ex post realized volatility. The various explanations for the existence of the forecast bias fall under one of two categories: either the options market is inefficient for some reason⁶, or, the option pricing model is incorrect. With regard to the latter, while the empirical literature treats implied volatility as being exogenous, a few researchers have suggested that there is an element of endogeneity in prices related to model misspecification. Option pricing models used in practice, such as the Black-Scholes model, do not allow for a premium for bearing volatility risk. The crux of the argument is that volatility risk premium exists, and so prices derived by models that do not incorporate this important variable will automatically be biased. By applying the Heston (1993) model, researchers have attempted to quantify the volatility risk premium as a way to assess the extent to which the

⁵For surveys of such time series modeling techniques that use empirical data, see Campbell, Lo and MacKinlay (1997), Gouriéroux and Jasiak (2001), and Taylor (2005).

⁶Some common explanations include the presence of trading frictions in hedging markets, liquidity frictions, and 'peso problems'.

forecast bias is due to its omission. Benzoni (2001), Chernov (2001), and Mixon (2007) all find that forecast bias is reduced by incorporating this dimension.

Furthermore, the bias will persist only if it is difficult or impossible to construct and efficiently implement arbitrage strategies that will in time remove the market mispricing. Fleming (1998) finds material forecast bias in option prices on the S&P 100 index. Arbitrage strategies designed to profit from mispriced equity index options would involve actively managing the underlying basket of stocks in the index, an administratively demanding process. Alternatively, the arbitrage may be done indirectly via index futures, which trade with great liquidity. In both cases, the arbitrage requires replication of the options, which involves frequent trading and active management of the position. This will discourage some of the potential arbitrageurs, thus contributing to the persistence of the bias. Measurement of realized volatility is also a potential source of bias persistence. Poteshman (2000) finds that a more efficient volatility estimator based on intra-day five minute returns removed over half of the bias present using daily data. Blair, Poon, and Taylor (2001) report up to a four-fold increase in r -squared coefficients when going from daily to high-frequency intra-day data.

2.1.B Forecast bias of the term structure of implied volatility

The second question addressed in the volatility forecasting literature, the one addressed in this paper, investigates the evolution of implied volatility itself across time. The existence of a term structure implies that the market places a different level of uncertainty about asset prices for different time horizons. In practice, this results in upward or downward sloping term structures, seldom flat. Under the benchmark Black Scholes (BS) model, the concept of a term structure should not exist. The assumption is that implied volatility is fixed across option tenors. The divergence between theory and real-world application has been the motivation for vast research efforts involving finance theorists, behavioral economists, and non-academic researchers seeking trading profit opportunities.

The forecast accuracy of implied volatility term structure has been tested in the literature by application of the expectations hypothesis. The intuition and tests are the same as those that have been applied to the term structure of interest rates, as per Shiller (1979) or Fama (1984). The theory suggests that the shape and slope of today's implied volatility term structure should link the long-end today with expected future changes of the short-end of the curve, or equivalently that the short-end today should be consistent with expected future changes of the long-end of the curve. Results when applied to financial options are mixed. Stein (1989) documents overreactions of the longer-dated option prices on the SPX index to changes

in short-dated options. Campa and Chang (1995) develop a well accepted framework for testing EH, and their results uphold the theory for a narrow set of currency option implied volatilities. Poteshman (2001) and Byoun et al. (2003) find the slope of the volatility term structure to have significant predictive ability for future implied volatility of the SPX. Mixon (2007) finds that by adjusting the implied volatility forecast by a volatility risk premium, the predictability along the term structure increases, however, not to the extent predicted by expectations hypothesis.

2.2 Volatility products

Volatility products, also known as volatility derivatives, have emerged over the last few decades due to the market's acceptance of volatility as being an asset class. The traditional role of volatility as a singular input to standard option products is upgraded for this class of securities. For a vanilla option, volatility plays a supporting role in the valuation of contingent non-linear payouts. Its contribution to option value peaks for at-the-money (ATM) strikes but falls very quickly as the moneyness of the option deviates higher or lower⁷. Expressing a view on future volatility via standard options is thus an inefficient and ineffective proposition, as the contribution of volatility to the value of the derivative is heavily contingent on the level or return of the underlying asset. Volatility products, in contrast, grant the holder a well-defined exposure to volatility itself. There are two basic types, written primarily on major equity indices or currency prices⁸. On one hand, there are variance (volatility) swaps whose payoff is linked to the observed realized variance (volatility) over the contract period. Such instruments may be thought of as forward contracts on future realized variance (volatility). On the other, there are forward volatility agreements, which pay off according to changes in the implied volatility itself⁹. These are essentially forward contracts on future implied volatility. Carr and Lee (2009) provide an excellent and thorough review of volatility derivatives.

2.3 VIX and VIX futures

The VIX is the most widely followed and well-known volatility product. Because the VIX is not an investable asset and because replication is costly, exposure to the index is achieved principally by trading VIX futures. The information content of futures markets is thus an important component for empirical work involving

⁷The vega exposure, or price sensitivity to changes in volatility, of a standard vanilla option decreases as the strike of the option goes further in-the-money (ITM) or out-of-the-money (OTM).

⁸The underlying instrument must trade with ample liquidity in order for construction of volatility derivatives to be possible.

⁹Implied volatility represents the market's estimate today about future realized volatility.

the VIX.

The VIX and the VIX futures market have been well covered in the literature across three dimensions: pricing, risk management, and forecast accuracy.

2.3.A Pricing

The construction and pricing of the VIX and other volatility products is based on the development of model-free implied volatility measures. At the fundamental level, model-free implied volatility can be derived from a portfolio of plain vanilla options with strikes spanning a wide range of possible outcomes for the underlying asset. The individual option prices in this approach are thus taken as given. Early work by Breeden and Lizenberg (1978) provided the foundation for describing theoretical relationships between underlying option prices and model-free implied volatility.

A separate strand of literature running in parallel involves the use of alternative processes for asset prices, volatility dynamics or alternative option pricing models to price volatility products. This approach does not accept traded option prices on the SPX as given. Zhang and Zhu (2006), for instance, use the Heston stochastic volatility model to price VIX futures. Zhu and Zhang (2007) value VIX futures by applying a stochastic variance model to the evolution of the VIX itself and to deriving the term structure of forward variance. Lin (2007) uses an affine jump-diffusion model with jumps in both index and volatility processes to arrive at VIX futures prices. Sepp (2008b) applies a similar framework for calibration of both VIX futures and options on the VIX. Zhang and Huang (2010) highlight the importance of dynamics assumptions and parameter estimation by contrasting the results of a number of different approaches.

2.3.B Risk management

Although not widely considered a stylistic fact, VIX-products are best utilized for purposes of diversification, and not to serve as true hedging vehicles. The difference is subtle but central to portfolio construction. Diversification involves decreasing portfolio variance, and thus portfolio risk, by adding elements that exhibit a low correlation to existing holdings. A hedge, on the other hand, is intended to offset losses to a core asset position (by exhibiting a high correlation). This is can be said of futures and options on the SPX, for instance, which derive their value directly from SPX price changes.

This key distinction was made by Daigler and Rossi (2006) who report a significant diversification benefit comes from adding a long VIX position to an SPX portfolio. Szado (2009) evaluates the performance

of adding a VIX-replicating portfolio to a base holding of stocks and bonds and finds a net reduction in aggregate risk of roughly one third. Delisle et al. (2010) go a step further and address the hedge inefficiency of VIX products by offering an alternative VIX-replicating portfolio which outperforms a static buy and hold position in the VIX index (assuming perfect replication is possible).

In Appendix A, I sketch out the basic framework for establishing VIX-holdings as being diversification vehicles, as opposed to true risk management hedges.

2.3.C Forecasting and the VIX

Because the VIX, by construction, is derived from a collection of all out-of-the-money options on the S&P 500, its forecasting accuracy should, conceptually-speaking, be better than implied volatility estimates extracted from a single option¹⁰. The literature is mixed on this front. Ait-Sahalia, Karaman, and Mancini (2012) address this question and find that the expectations hypothesis, defined as the difference between forecasted and actual variance, does not hold, but biases and inefficiencies are modest for short time horizons. Becker, Clements and White (2006) examine whether the VIX contains any information relevant to future realized volatility beyond that reflected in model-based estimates based on empirical data. The authors conclude that the VIX does not add to the forecasting power of alternative approaches. Shortly after, Becker and Clements (2008) show that a combination of both the VIX and a model-based estimate is found to be superior to either method by itself.

In appendix B, I assess the forecast accuracy of the VIX itself according to its strict definition. I find that ex ante estimates given by the VIX index overshoot ex post realized volatility in the SPX. These results are consistent with the prior academic work on this topic. Option prices are biased upward, as compared to subsequent payouts.

2.3.D Forecasting and the VIX futures market

A number of different studies involving forecast accuracy have been applied to VIX futures markets. Konstantinidi, Skiadopoulou, and Tzagkaraki (2008) and Konstantinidi and Skiadopoulou (2011) demonstrate that VIX futures are predictable by their historical patterns, however the coefficients are too small to generate actual trading profits. Bollerslev, Tauchen, and Zhou (2009) and Drechsler and Yaron (2011) study the relation between variance risk premia and equity risk premia and demonstrate that the VIX has significant

¹⁰The literature on options efficiency has primarily focused on implied volatility estimates based on individual options, at-the-money strikes are most common.

forecasting ability for future SPX returns. Their models are based on the underlying VIX, determining that no additional information is being furnished by the term structure. Lu and Zhu (2010) study term structure dynamics and identify that a third factor, capturing curvature movements, is statistically significant. Shu and Zhang (2012) apply traditional linear Granger tests and find that VIX futures prices have predictive ability on the underlying VIX, however after exploring for non-linear relationships, the predictive ability goes away. Luo and Zhang (2012) propose two-factor stochastic volatility framework for the VIX, corresponding to the level and slope of the VIX term structure, which offers rich information content relative to historical volatility in forecasting future realized volatility.

Regarding tests of EH, Nossman and Wilhelmsson (2008) operate under a stochastic volatility framework in order to adjust actual VIX futures prices by a risk premium and find that EH cannot be rejected. Also they report that risk premium adjusted futures prices provide good forecasting ability of the VIX index with a 73% hit ratio. A potential limitation in their work however is that the researchers only use the near-term VIX futures contract with a maximum tenor of thirty days. In addition, their calibration of risk premium grows linearly with time to maturity, with mean and upper bounds at two and four percent respectively for a thirty-day period. Extrapolating this to a six-month contract, which trades with ample volumes, would deem this an unrealistic metric. More recently, Ait-Sahalia, Karaman, and Mancini (2012) find that a significant jump risk component embedded in the VIX term structure contributes to a downward sloping term structure in quiet times, but an upward slope in turbulent times.

3. VIX defined

3.1 Index construction

The VIX represents the conditional risk-neutral expectation of the volatility for the SPX index over the next calendar month

$$\sigma_t^{VIX} \cong E_t^0[\sigma_{SPX_{t \rightarrow t+1}}], \quad (1)$$

where σ^{VIX} is the estimate of volatility as quoted in annualized percentage terms according to standard market convention, and σ_{SPX} represents the realized volatility of the SPX, also expressed in annualized percentage terms, from time t to $t + 1$ months later. The methodology for arriving at VIX prices involves computing a weighted average of out-of-the-money (OTM) option prices on the SPX across all strikes for the

two nearby maturities¹¹. The general formula used in the VIX calculation¹² is given by (2) and (3) below

$$\sigma_{p_j}^2 = \frac{2}{T_j} \sum_i \frac{\Delta K_i}{K_i^2} e^{rT} Q(K_i) - \frac{1}{T_j} \left[\frac{F_j}{K_0} - 1 \right]^2, \quad (2)$$

where $\sigma_{p_j}^2$ is the constructed variance of a portfolio of SPX options expiring at the two nearby maturities j and $j + 1$, T , expressed in years, is the common period over which all options in the calculation are active, F is the forward index level derived from coterminal index option prices¹³, K_0 is the first strike below the forward, K_i is the strike price of the i^{th} out of the money option¹⁴, ΔK_i denotes the interval between strike prices, defined as $\Delta K_i = (K_{i+1} - K_i)/2$, r is the risk-free rate to expiration, $Q(K_i)$ is the midpoint between the bid-ask spread for each option with strike K_i , and r is the risk-free rate.

Table 1 contains a collection of options on SPX that would be involved in a hypothetical calculation of the VIX as of 2-Nov-2012. In practice, option strikes are available for every five points on the SPX index, however for expositional purposes, the intervals used are fifteen and twenty-five SPX points apart for near and next-term maturities respectively.

- Insert Table 1 here -

The salient information to be extracted from the table is as follows: 1) The strike of 1415 is the strike where the price difference between calls and puts for both maturities is smallest, 2) the strike of 1415 will be used to determine F and K_0 , which in turn determines the set of K_i , 3) the range of option strikes used for each maturity will vary as the calculation leaves out options for which the bid price is zero, and 4) the exact collection of options used will change in tandem with changes in the underlying price of the SPX as in-the-money (ITM) options are left out of the calculation.

The VIX, as quoted, is computed by deriving (2) for the near-term and the next-term maturities on SPX futures, indexed by j and $j + 1$ respectively¹⁵, and plugging below

$$\sigma_t^{VIX} = \sqrt{\left\{ T_j \sigma_{p_j}^2 \left[\frac{N_{T_{j+1}} - N_{30}}{N_{T_{j+1}} - N_{T_j}} \right] + T_{j+1} \sigma_{p_{j+1}}^2 \left[\frac{N_{30} - N_{T_j}}{N_{T_{j+1}} - N_{T_j}} \right] \right\} \frac{N_{365}}{N_{30}}}, \quad (3)$$

where N is the number of minutes for each referenced period. As suggested by the weighted average calcu-

¹¹Maturities occur monthly. In eight out of twelve months in the year, VIX futures settle on the third Wednesday of each month, in the other four months, the futures expire on the fourth Wednesday of the month.

¹²According to white paper on the VIX accessible on the CBOE website (www.cboe.com/micro/vix/vixwhite.pdf).

¹³This is equivalent to the strike price at which the price difference between the SPX call and put is smallest adjusted by e^{rT} (call price - put price).

¹⁴Note if a call option, then $K_i > K_0$ and if a put option, $K_i < K_0$.

¹⁵The near-term future is the 1st future, and the next-term future is the 2nd future. These terms will be used interchangeably.

lation in (3), VIX prices are updated continuously throughout the day to the very minute. Figure 2 shows the output of this calculation on an intra-day basis.

- Insert Figure 2 here -

Although not perfect, the strength of the co-movement is evident.

3.2 VIX futures

The VIX is an index, not an investable asset. The core method for attaining exposure to the VIX is via VIX futures, which began trading on the CBOE on 26-March-2004. VIX futures are essentially forward contracts on future implied volatility. There is however no cost of carry relationship between the VIX and VIX futures, as is standard between spot and futures prices of other exchange-traded assets¹⁶. Instead, by way of its construction, a position in VIX futures is an expression which links today's expected volatility to tomorrow's expected volatility,

$$\mathcal{F}_{t,j \rightarrow j+1}^{VIX} = E_t^{\mathbb{Q}}[\sigma_{j,j \rightarrow j+1}^{VIX}], \quad (4)$$

where \mathcal{F} represents the conditional risk neutral expectation at time t of the VIX at the first future date of j . Note the VIX is always a one-month forecast of realized volatility, thus \mathcal{F} is a forecast of this one-month forecast. The tenor of \mathcal{F} , or the length of the forecasting period, will vary however. Combining equations (1) and (4) allows us to express futures prices as

$$\mathcal{F}_{t,j \rightarrow j+1}^{VIX} = E_t^{\mathbb{Q}}[E_j^{\mathbb{Q}}[\sigma_{SPX_{j \rightarrow j+1}}]], \quad (5)$$

the expectation at time t , of the expectation at the time of the first future date j , of the realized volatility of the SPX index over the period j to $j + 1$. Figure 3 illustrates the interaction between these concepts.

- Insert Figure 3 here -

On 1-February-2012, time t , the closing price quotes, in annualized standard deviation terms, for the VIX and the 1st and second futures were $\sigma_t^{VIX} = 18.55$, $\mathcal{F}_{t,j \rightarrow j+1}^{VIX} = 19.85$, and $\mathcal{F}_{t,j+1 \rightarrow j+2}^{VIX} = 22.05$ respectively.

¹⁶Similarly, VIX futures prices do not contain elements related to insurance, storage, and transportation costs. This represents a departure from the extensive literature that deals with the forecast accuracy of futures prices for other assets such as commodities, currencies, interest rates.

Note that σ_t^{VIX} is a forecast of realized volatility of the SPX over the immediate future period, while $\mathcal{F}_{t,j \rightarrow j+1}^{VIX}$ and $\mathcal{F}_{t,j+1 \rightarrow j+2}^{VIX}$ are forecasts of the future forecasts of realized volatility of the SPX. The horizontal dashed arrows in figure 3 represent the period over which such expectations apply.

As conveyed previously, the focus of this paper is not to explore the extent to which the VIX accurately forecasts the future realized volatility of the SPX. Rather, the objective is to understand and characterize the evolution of the expectations. For illustrative purposes, a similar distinction can be made within the insurance industry. Expectations of floods in New York City is best captured by evaluating the price level and changes of flood insurance premiums. An analysis of the historical revenues earned from these premiums measured against subsequent payouts is also important, and indirectly contributes to the actuarial fair value of premiums, however, it does not capture the dynamic nature of the expectations of a flood on the part of end-users such as businesses and households.

3.3 Pricing VIX futures

While VIX prices are derived from an exact calculation, VIX futures prices are ultimately determined by supply and demand dynamics in the market. The actuarial fair value a VIX futures contract can be determined from a synthetic calendar spread¹⁷ of SPX options bracketing the one-month period which starts at the futures expiration date, minus a term which estimates the risk-neutral variance estimate of VIX future. The derivation for (6) can be found in Carr and Wu (2006). Using the example in the previous section, the fair value of the next-term future is thus given by

$$\mathcal{F}_{t,j+1 \rightarrow j+2}^{VIX} = \sqrt{[\mathbb{P}_t - \hat{\sigma}_{\mathcal{F}_{t \rightarrow j+2}}^2]}, \quad (6)$$

where \mathbb{P} represents a portfolio of SPX options with long positions in out-of-the-money options expiring in $j + 2$ and short positions in out of-the-money options expiring in $j + 1$. The second term, $\hat{\sigma}_{\mathcal{F}_{t \rightarrow j+2}}$, denotes an estimate of the cumulative variance of \mathcal{F} between t and $j + 2$, thus it is not known at time t . It represents the concavity adjustment required since the static portfolio construction per equation (2) characterizes the variance profile, and does not sufficiently capture volatility, the square root of variance. In simple terms, if the VIX was quoted as variances and not annualized volatilities, this adjustment would not be required. Of course, expressing the VIX in volatility terms naturally increases the marketing appeal of the index, as

¹⁷A calendar spread is a trade involving the simultaneous sale and purchase of a pair of futures or vanilla options expiring on different dates. The legs of the spread vary only by maturity, as they are based on the same asset and notional amounts. The rationale for entering this trade is to take advantage of perceived value along the term structure, to partially finance long positions, or to cap the maximum loss of short positions.

volatility is easier to understand and interpret by market professionals, retail investors, and researchers¹⁸.

As mentioned, the term $\hat{\sigma}_{\mathcal{F}_{t \rightarrow j+2}}$, by definition an estimate of the variance of the underlying futures price, can never be known exactly at time t as it covers the period from t to expiration of the futures contract. Researchers have aimed to address this issue by applying different assumptions about asset and volatility dynamics in their construction of synthetic VIX futures curves using equations (2) and (3)¹⁹. Indirectly such calibrations, in one way or another, aim to estimate this unknown and dynamic parameter.

In this paper, I will use actual VIX futures prices to construct and analyze the term structure. Thus, $\hat{\sigma}_{\mathcal{F}_{t \rightarrow j+2}}$, is embedded in the price quote. This represents a contribution to the field as prior studies construct the term structure using the formulation furnished by the CBOE, described in (2) and (3). Futures prices offer economically relevant information about the market's expectation of the variance profile of the futures contract throughout its life. In fact, the shorter the tenor, the more its pricing is impacted by this estimate of unrealized future variance, in relative terms. Another stylistic fact about volatility as reported in Poon and Granger (2003) is that forecasts of cumulative volatility become more accurate as the period of time over which volatility is generated grows due to cancellation of error and mean reversion dynamics. Application to VIX futures fair value pricing per equation (6) suggests

$$Var(\hat{\sigma}_{\mathcal{F}_{t \rightarrow j+6}}) < Var(\hat{\sigma}_{\mathcal{F}_{t \rightarrow j+3}}) < Var(\hat{\sigma}_{\mathcal{F}_{t \rightarrow j+1}}) . \quad (7)$$

As the tenor of the futures contract decreases (increases), the uncertainty around the concavity adjustment rises (falls). This has two key implications. The first is that all else equal, shorter-dated futures should exhibit greater actual variance than longer-dated futures. This is a seemingly circular argument, but the intuition is simple. If the fair value pricing today is a function of projections about tomorrow's variance, and there is greater potential variability about this variance, then the fair value itself should exhibit greater variance. This is confirmed by the descriptive statistics in Table 2 in a later section. The second implication is that an options market that develops on the actual VIX futures themselves should be characterized by downward sloping implied volatility curves. This is also the case, although options on VIX futures are not covered here.

¹⁸There is such a product that is written on variances, but it is not very widely traded for reasons discussed. The S&P 500 3-month variance futures, which began trading on the CBOE on 18-May-2004, pay off according to the calculation of realized volatility over the period. The 12-month variance future, also launched in 2004, was delisted in 2011. Plans for a re-launch date of 4-October-2012 were postponed.

¹⁹As cited in section 2, see Zhang and Zhu (2006), Zhu and Zhang (2007), Lin (2007), Sepp (2008b), etc. for examples.

4. The framework for evaluating forecast bias

The expectations hypothesis will be applied to evaluate the forecast accuracy of VIX futures, that is, the extent to which the evolution of the VIX ex post is described by VIX futures price quotes ex ante along various tenors. The focus is to evaluate changes in expectations of future volatility, not expectations versus actual realized volatility which is a separate question.

4.1 Theory

Campa and Chang (1995) derive the relation for testing EH under three key assumptions: 1) volatility is stochastic, 2) the underlying asset and its volatility are uncorrelated, and 3) there is no volatility risk premium. The first assumption is a departure from the Black Scholes framework, but a well accepted concept within option pricing theory. As Hull and White (1987) demonstrate, the discrepancy between option prices using stochastic versus static volatility is independent of the level of volatility but increases with the time to maturity. The second assumption is a strong one, and generally not true in practice. A fall in asset prices is generally associated with a rise in volatilities, as the demand for protection rises. On the other hand, rising asset prices results in less demand for protection and thus a decrease in implied volatilities. This is evident across equities, currencies, and other asset classes. Mixon (2007) tests EH using a similar formulation and points out that nonzero correlation would not present a problem provided there is no material wedge between the average expected volatility and the implied volatility for ATM options. Finally the last assumption is that volatility risk premium is zero. Campa and Chang do not test this assumption directly but instead present empirical evidence which suggests the assumption is plausible. They determine that even if a risk premium exists but not considered, the bias from omitted variables in their specification would contribute to further support of EH. Mixon (2007) expands on this work by defining volatility dynamics per Heston (1993) and proposes a formulation that is inclusive of a risk premium, thereby departing from Campa and Chang's risk neutral framework to an objective measure as suggested by the data. Although theoretically sound, his approach relies on estimates of future realized volatility which he furnishes using a GARCH (1,1) specification²⁰.

In this paper, I will utilize Campa and Chang's risk neutral formulation to test EH. The rationale is as follows. For one, the VIX by construction does not constrain volatility to be constant. The index pools option prices for all the out-of-the-money strikes, thereby extracting all information conveyed by the short-

²⁰Generalized Autoregressive Conditional Heteroskedasticity

end of the term structure and the implied volatility skew²¹ per equations (2) and (3). With regard to the second assumption, I evaluate the implied volatility skew of options on the VIX index itself and based on Mixon (2007) confirm that there is no material wedge between the ATM strike, which is essentially the spot VIX, versus the average implied volatility of all ITM and OTM strikes for both calls and puts. Regarding the third assumption, as in Campa and Chang (1995), the omitted variable bias from using a specification that leaves out risk premiums would provide further support for EH. In other words, the regression coefficients are closer to unity, which contributes to forecast accuracy.

4.2 Constructing the VIX term structure

The innovation proposed in this paper involves testing EH based on the construction of the VIX term structure, which are essentially hypothetical variance swaps²², using actual VIX futures market prices. Nossman and Wilhelmsson (2008) test EH on actual VIX futures, but use only the near-term contract. As outlined previously, researchers have estimated the VIX term structure according to a model-free approach by replicating the CBOE formulas (2) and (3), many times under a wide variety of assumptions about asset and volatility dynamics, as well as alternative models for option pricing. Ait-Sahalia, Karaman, and Mancini (2012) generate the VIX term structure in the presence of jumps for the underlying asset for instance.

Whereas the standard VIX represents the risk neutral expectation of the realized volatility of the SPX over a one-month period per equation (1), the VIX term structure represents long-dated VIX contracts which are estimates of the realized volatility over a period of more than one month. For example, the long-dated VIX contract which expires at the next-term or second VIX future date, $j + 1$, is defined as

$$\sigma_{t,t \rightarrow j+1}^{VIX} \cong E_t^Q[\sigma_{SPX_{t \rightarrow j+1}}]. \quad (8)$$

Then applying the definition of expectations hypothesis allows constructing the variance swap which expires at the next term date as the sum of the VIX plus the near-term VIX future

$$E_t \left[T_{j+1} \left(\sigma_{t,t \rightarrow j+1}^{VIX} \right)^2 \right] = E_t \left[T_j \left(\sigma_{t,t \rightarrow j}^{VIX} \right)^2 \right] + E_t \left[(T_{j+1} - T_j) \left(\mathcal{F}_{t,j \rightarrow j+1}^{VIX} \right)^2 \right], \quad (9)$$

where T is the time period expressed in years ending at the time specified by the subscript j , j represents

²¹Term structure refers to implied volatility differences for different maturities, while the skew refers to implied volatility differences for different strikes for a given tenor.

²²The terms variance swap and long-dated VIX contract will be used interchangeably, both represent hypothetical instruments that extend VIX beyond a one month tenor. The VIX at the various future tenors constitutes its term structure. This is not the same concept as the VIX futures curve.

the near-term or 1st future expiry date, $j + 1$ represents the next-term future expiry date, $j \rightarrow j + 1$ implies that the future fixes at j but yields a forecast of realized volatility on the SPX to $j + 1$, and t denotes today. The left hand side of (10) constitutes the long-dated VIX contract or variance swap, while the right hand side contains elements that are available per market pricing of the VIX and VIX futures. It is important to note that the calculation for the long-dated VIX contract which expires at the next future date does not contain the VIX future that expires at the next future date. It contains only the VIX and the VIX future that expires at the near-future date. Figure 3 illustrates the intuition behind this one period lag. I proceed this way and construct VIX term structure out seven months. The VIX as quoted represents the first point on the VIX term structure.

Figure 4 shows the resulting VIX term structure based on the sample quotes from an earlier section.

- Insert Figure 4 here -

The constructed VIX term structure will be lower than the VIX futures curve when the slope of the VIX futures curve is positive. Forecasts of cumulative volatility become more accurate as the period of time over which volatility is generated grows due to cancellation of error and mean reversion dynamics. Thus a six-month long-dated VIX contract or variance swap represents the expected variance in the SPX over a six-month period, while the six-month VIX future represents the estimate today for the VIX index, or equivalently the thirty-day estimate of the variance in the SPX, six months from now. The latter naturally carries greater uncertainty, and thus a higher volatility price.

4.3 The testable equation for expectations hypothesis

Using the results of Hull and White (1987), the derivation by Campa and Chang of the testable equation for EH starts with establishing that the appropriate price for an option quoted at time $t = 0$ equals the Black-Scholes price, evaluated at the average variance over the life of the option

$$C^{HW} = E[C^{BS}(\bar{\sigma}_{0,t}^2)], \quad (10)$$

where C^{HW} and C^{BS} denote the prices under Hull and White and Black-Scholes respectively. It is well known, however, that the Black-Scholes model evaluated at the average variance overprices ATM options, and this overpricing increases as the time to maturity grows. And thus, for an ATM option, the concavity in σ implies

$$C^{HW} = \theta_t C^{BS} E[(\bar{\sigma}_{0,t}^2)], \quad (11)$$

where θ_t corrects for the mispricing, is less than one, decreases as time to maturity grows. The concavity is small for ATM options, and thus (12) can be expressed as

$$\theta_t C^{BS} E[(\bar{\sigma}_{0,t}^2)] \cong C^{BS} \left[\theta_t^2 E(\bar{\sigma}_{0,t}^2) \right], \quad (12)$$

relating the average variance over the life of the option to the implied Black-Scholes variance. Applying the law of iterated expectations on the definition of EH, Campa and Chang show the relation for current and future expected ATM implied volatilities, expressed in variances or volatility squared

$$\sigma_{0,km}^2 = \left(\frac{1}{k} \right) E_0 \left[\sum_{i=0}^{k-1} \sigma_{i,m,(i+1)m}^2 \right] \left(\frac{\theta_{km}}{\theta_m} \right)^2, \quad (13)$$

where m is the number of months until expiration for the short-dated option, k as the number of periods of length m , and θ is the concavity adjustment for a given tenor. Equation (13) says that the current volatility quote equals the average of the current and expected future short-dated volatility quotes. In other words, the slope of the term structure is informative about where implied volatility will be in the future. Next, I apply the same simplifying assumption, $\theta_{km}/\theta_m = 1$, as in Campa and Chang (1995). There are pros and cons to this avenue. The benefit is that the testable equation becomes more streamlined and intuitive. The cons are that the omission of terms results in the introduction of omitted variable bias. If the covariances between the term structure level and the term structure slope are negative (positive), then the omitted variable bias results in regression coefficients (for tests of expectations hypothesis) that are biased downward (upward). A downward bias for coefficients is acceptable, while an upward bias is not. Expectations hypothesis is upheld in this context if coefficients are larger (and closer to 1.0). Thus failure to reject the hypothesis test using a specification that is biased downward is conservative, but if the specification is upwardly biased, then this may be a false positive. I proceed with this specification, as the covariances aforementioned are in fact negative²³.

Since the level of volatility follows a near unit-root process, equation (14) is tested in terms of long-short spread rather than using variance levels directly. Then, subtracting the current short-dated option variance, $(\sigma^{VIX})^2$, to both sides yields the testable equation for the expectations hypothesis on the VIX

²³The time series of the seven-month point on the term structure is negatively correlated to the term structure slope between one and seven months, and so on respectively for the other points on the term structure.

term structure. Returning to the notation used in this paper yields

$$\left(\frac{1}{k}\right) \sum_{i=1}^{k-1} \left[(\sigma_{i,i \rightarrow i+j}^{VIX})^2 - (\sigma_{0,0 \rightarrow j}^{VIX})^2 \right] = \alpha_0 + \beta_0 \left[(\sigma_{0,0 \rightarrow k+j}^{VIX})^2 - (\sigma_{0,0 \rightarrow j}^{VIX})^2 \right] + \sum_{i=1}^{k-1} u_i, \quad (14)$$

where u_i represent the expectational errors. It should be noted from equation (15) that the VIX futures as quoted, \mathcal{F} , enter the regression as the constructed long-dated VIX contracts, or variance swaps, per the calculation in equation (9).

5. Data

The dataset includes daily closing prices for the SPX, the VIX, and VIX futures from January 2006 to November 2012. The original source of this data is the CBOE, however, they were gathered through Bloomberg. In addition, daily price and volumes for ETF offerings linked to the VIX will also be used.

5.1 VIX and VIX futures

A single reading per day is available for the VIX index. The frequency for the analysis will be weekly, however. Weekly data is selected as it helps reduce the forecast bias introduced by including overlapping forecasts in the analysis, which produces serially correlated errors. In addition, each day there is a strip of VIX futures contracts that expires at set dates in the future. The near or 1st future expires in the same month, the next or 2nd future expires in the following calendar month and so on. When VIX futures were launched on 26-March-2004, there were a total of four futures contracts trading on any given day. Today there are total of nine futures contracts, with the longest maturity being more than one year out. In order to utilize as much history as possible in this analysis, I include the first six months of maturities only, and focus on tenors two, three, and six. This allows merging the older data with the newer data, without having to rely too much on interpolation and extrapolation to fill missing maturities. Table 2 contains descriptive statistics on the VIX, VIX futures, and the resulting VIX term structure spanning the period 2006-2012.

- Insert Table 2 here -

The average figures suggest there is a tendency for the VIX futures curve to be upward sloping. In addition, the ranges of prices for the VIX as well as the various points along the VIX futures curve and

VIX term structures convey that the volatility of the prices themselves decreases as the time to maturity increases, and this is confirmed by the standard deviation readings. The intuition for this was discussed in the VIX futures pricing section. Essentially as the tenor of the futures contract rises, the uncertainty around the concavity adjustment should fall. Skewness readings are in line with implied volatilities in other asset classes. Kurtosis figures indicate fat tails for the underlying VIX, but not for the set of futures. The short-end of the resulting VIX term structure displays fat tails.

Figure 5 tracks the spread between six-month VIX futures (left axis) and the underlying VIX (right axis) for the entire sample period.

- Insert Figure 5 here -

A few observations worth noting. For one, roughly 75% of the time, the spread between the long-end future and the VIX is positive. Also, the spread went steeply negative during the height of the 2008 credit crisis and then materially negative in the latter half of 2011 when Euro zone crisis fears peaked. And three, the average spread when positive was substantially greater post-crisis, as opposed to pre-crisis.

5.2 About the curves associated with the VIX

Figure 6 is a snapshot of the various curves associated with the VIX as of the close of business on 3-November-2010, the day the second round of quantitative easing (QE2) was announced in the United States.

- Insert Figure 6 here -

The VIX index closed at 19.56, while the 6-month VIX future closed above 25, displaying a significant premium. As discussed, when the VIX futures curve is upward sloping, the constructed VIX term structure or variance swap curve, will generally be lower. Also depicted in the figure are the term structures for at-the-money (ATM) and 25-delta out-of-the-money (OTM) implied volatilities for SPX options. SPX options are the building blocks of the VIX index. It is common for the term structure of ATM volatilities on underlying SPX options to be lower than the VIX term structure across tenors, as the latter prices in the skew that is prevalent for OTM options on all financial assets. By construction, the VIX gives the holder exposure to the full set of OTM options on SPX at any one particular time, and thus the skew is reflected in the higher term structure.

Regarding the relationship between the VIX term structure and the term structure of 25-delta OTM SPX options, we would expect the latter to be higher on average. The payoff of a variance swap such as the VIX is convex in volatility. This means that an investor who is long a variance swap will benefit from boosted gains and discounted losses, a phenomenon which is amplified when volatility skew is steep. Thus, the fair strike of a variance swap is often in line with the implied volatility of 40-delta SPX puts, which carries a lower implied volatility than that of a 25-delta put in the presence of skew adjustments.

6. Testing the forecast bias of the VIX term structure

Expectations hypothesis will be tested to address the following research question. When the VIX term structure is positively (negatively) sloped, does the VIX subsequently rise (fall) as much as predicted? Specifically, equation (14) evaluates the long-short spread of the constructed VIX term structure at time $t = 0$ against subsequent changes in the VIX over two, three, four, five, six and seven-month periods. The long-short spread is an unbiased predictor of the future movements in the VIX if $\alpha = 0$ and $\beta = 1$. This would imply that for every unit of variance in the long-short spread of today, we would expect one unit of variance in subsequent movements in the VIX index. Rejections of this joint hypothesis test will be reported at the 5% level.

In addition, while it is useful to have unbiased forecasts, biasness and predictive power are two separate concepts. As Poon and Granger (2003) point out, a biased forecast can have predictive power, but an unbiased forecast is useless if forecast errors are always big. For this reason, I also run a hypothesis test to evaluate the extent to which $\beta = 1$. Rejection of a hypothesis test based on this single coefficient would not resolve the issue of biasness, as a joint hypothesis test would be required, however the test is important from the standpoint of highlighting potential arbitrage opportunities. I also run a third test for investigating whether $\beta = 0$. Rejecting this hypothesis would confirm that the VIX term structure contains valuable information content. Otherwise, if $\beta = 0$, then changes in the VIX are white noise adjusted by a constant term.

6.1 Results

Table 3 contains the results for the joint hypothesis test $\alpha = 0$ and $\beta = 1$, and the individual tests $\beta = 1$ and $\beta = 0$. The regressions are based on weekly readings from January 2006 to February 2012. Weekly data

is selected as it helps reduce the forecast bias introduced by including overlapping forecasts in the analysis, which produces serially correlated errors. The serial correlation is also addressed by applying error correction techniques outlined in Newey and West (1987)²⁴.

- Insert Table 3 here -

The joint hypothesis test results dictate that the long-short spread does predict the direction of the subsequent move in the VIX correctly, since the beta coefficients are positive, but for the most part, not to the extent implied by the expectation hypothesis, since they are all less than one. The coefficient reading of 0.802 for the seven-month point on the VIX term structure implies that ten units in today's long-short spread are associated with a rise in the VIX index of 8.02 units over a seven-month period. The beta coefficients for all points on the curve, with the exception of the 7-month, are all significantly less than one at the 5% level, implying that VIX futures are consistently overpriced. The forecast bias increases substantially with shorter tenors. This is a feature also prevalent in the results for Campa and Chang (1995) and Mixon (2007), confirming a well-known stylistic fact about option-pricing. Long-dated options offer better value to hedgers and speculators as opposed to short-dated options. The results of the second hypothesis test, which evaluates whether $\beta = 1$ individually, does not reveal any new information over the joint test. The results of the third hypothesis test, which evaluates whether $\beta = 0$, establishes that there is no information content in the VIX term structure over the short-part of the curve only (at two and three-month tenors).

6.2 Sub-period analysis

To control for the time period effect, I run the hypothesis tests for three disjoint periods surrounding the 2008-9 credit crisis period: pre, during, and post. Table 4 contains the results for the joint hypothesis test $\alpha = 0$ and $\beta = 1$, and the individual tests $\beta = 1$ and $\beta = 0$ for the pre-crisis sub-period.

- Insert Table 4 here -

With exception of the seven-month point on the curve, the beta coefficients are closer to one in comparison to those of the entire sample, and the improvement is greatest at the short-end. In addition, the standard errors are larger in this sub-period, contributing to joint hypothesis test results which suggest no forecast

²⁴The lags for this error correction technique will be selected according to the point on the term structure being evaluated. The longer the tenor, the greater the lag.

bias for the four through seven month points on the curve. During this sub-period, the VIX averaged 12.8%, the six-month long-short spread of the VIX term structure averaged 1.6%, and the six-month VIX future averaged 2.4% above the underlying VIX index. The term structure was consistently forecasting an appreciation in the VIX, and the forecast was upheld for the term structure at months four, five, six and seven.

Table 5 contains the results for the joint hypothesis test $\alpha = 0$ and $\beta = 1$, and the individual tests $\beta = 1$ and $\beta = 0$ for the sub-period that covers the credit crisis of 2008-9.

- Insert Table 5 here -

This was a highly volatile period. Financial option prices rose across asset classes, commensurate with increased market uncertainty. The average level of the VIX index was 31%, more than double the pre-crisis period, and peaked at over 70% following the Lehman bankruptcy in the fourth quarter of 2008. Also the term structure was inverted for a majority of this period, a common feature that arises during times of financial distress. The average spread between the 6-month VIX future and the VIX was -3.9%. Results show beta coefficients are higher across the board, in fact, for the seven-month point on the term structure, it is above one, which suggests subsequent move in the VIX was greater than forecast by the term structure. The hypothesis test results demonstrate no forecast bias for the five, six, and seven points on the curve. The intuition from these results suggests that options are well-priced during periods of extreme volatility, despite the higher premiums paid.

Table 6 contains the results for the joint hypothesis test $\alpha = 0$ and $\beta = 1$, and the individual tests $\beta = 1$ and $\beta = 0$ for the sub-period following the credit crisis of 2008-9.

- Insert Table 6 here -

Post-crisis, there are rejections of the expectations hypothesis across the board. The beta coefficients are further away from one as compared to the other sub-periods, in fact, they are negative for the two and three-month points on the curve. A negative coefficient implies that a positive slope on the term structure was associated with subsequent falls in short-dated volatility forecasts (of realized volatility in the SPX). The average reading on the VIX was 24.3% for this period, roughly a quarter below the crisis period average. However, the slope of the term structure was consistently positive and steep. The average long-short spread of the VIX term structure at the six month point was 2.9% and the average spread between the 6-month

VIX future and the VIX was 3.8%. The divergence between current and future expectations is an important phenomenon. The overall level of risk as determined from the VIX was lower during the post-crisis period, but the expectations for the future level of the VIX to rise as extracted from the term structure slope were among the highest levels registered. The hypothesis tests determined that the expectations of subsequent rise in the VIX based on the positively-sloped term structures did not manifest.

To summarize the results of the tests, I will highlight the some important observations. First, VIX futures are consistently overpriced relative to the subsequent moves in the underlying VIX index. Second, the forecast bias increases substantially with shorter tenors. Third, the forecast bias is smallest during periods of extreme volatility. Four, deviations from expectations hypothesis are greatest during the post-crisis period.

7. What drives the forecast bias?

Now that the existence of the forecast bias of the VIX term structure has been documented, characterized and quantified, the focus shifts to addressing why it exists and why it persists. A number of possible factors might influence the size of the forecast bias over time. This section includes a discussion of each, and a description of the specific variables that will be used to establish and evaluate the linkage quantitatively in the next section.

7.1 Open interest

The VIX itself is not a tradeable asset. It is not possible to speculate or invest in the index directly. Exposure to the VIX is attained by synthetically replicating the index using portfolios of individual options on the S&P 500 according to the methodology outlined by the CBOE, trading futures and options whose payoff is determined by future levels of the VIX, or via ETF's or managed portfolios offered by the investment management industry. These methods are listed in order of decreasing difficulty of implementation.

Over the years following the credit crisis of 2008-9, the proliferation of ETF offerings involving the VIX has opened the flood gates and capital has poured into the strategy. Today an investor does not have to replicate the index or trade VIX futures or options on the VIX to get exposure to the index, both challenging avenues for individual investors, family offices, and smaller investment funds. Instead, an investor may buy or sell individual ETF shares, via a number of online portals, at very low transaction costs. VIX-related ETF's, along with managed funds offerings, have quickly grown to a market capitalization estimated at

\$3-5 billion between 2009 and 2012. The market capitalizations of three of the larger funds, according to Bloomberg data, are as follows: 1) the iPath S&P 500 VIX short-term futures fund with \$1.56 billion in assets (ticker VXX), 2) the iPath S&P 500 VIX medium term fund with \$116m in assets (ticker VXZ), and the iPath S&P 500 VIX dynamic fund with \$309m in assets (ticker XVZ), as of November 2012.

The transmission mechanism from ETF's to VIX futures markets is simple. In order to provide individual investors with a risk profile that tracks the underlying VIX index, professional money managers must trade futures and, to a lesser degree, options on the VIX. Mass buying of VIX ETF's then translates into large amounts of capital flowing into futures markets. A number of different portfolio offerings are available, based on varying strategies. The VXX ETF, the largest by market capitalization, offers exposure to the short-end of the term structure and thus a risk profile that most closely mimics that of the underlying VIX index²⁵. The XVZ ETF, in contrast, offers exposure to the entire term structure depending on perceived value at the discretion of the money manager. Figure 7 tracks the open interest in VIX futures along with the market capitalization of the VXX ETF.

- Insert Figure 7 here -

Open interest for VIX futures markets averaged approximately 50,000 contracts prior to 2009. After 2009, there was a sharp rise in open interest, no doubt related to the proliferation of VIX ETF offerings, as seen in figure 7. Today, ETF's represent the main vehicle for accessing the VIX, and inflows into ETF's implies inflows into the market for VIX futures. Open interest may play a key role in explaining the magnitude of the forecast bias in the VIX term structure, thus, I include these data in the regression. The source of open interest data is the CBOE, however, the data was gathered through Bloomberg.

7.2 Slow-moving capital

Application of the theory of slow-moving capital, as suggested by Gromb and Vayanos (2002) and Brunnermeier and Pedersen (2009) stresses that slow-moving capital may play a key role in propagating mispricing in financial markets. Is it plausible that changes in the availability of capital that could potentially be directed towards arbitrage strategies contributes to changes in the forecast bias in the VIX term structure? I follow Matthias, Longstaff, and Lustig (2013) in testing for this by including changes in total global hedge fund net

²⁵The descriptive statistics in table 2 demonstrate that the distribution of the short-end of the term structure more closely matches the distribution of the underlying VIX index, as compared to the longer-end of the curve.

asset values as one of the variables in the regression. The ticker of the specific time series from Bloomberg is HFRXGL Index²⁶, however, the original source of this data is Hedge Fund Research Inc..

7.3 Performance of safe haven assets

A number of investments have emerged as potential safe havens following the 2008-9 credit crisis: US treasuries, gold and other precious metals, agricultural commodities, the Japanese yen, the Swiss franc, the Chinese renminbi, land, real estate, etc.. Some have fared better than others²⁷, some are more accessible than others²⁸, while some have lost their safe haven status altogether²⁹. For the regression analysis, I will use the price of gold as representative of this asset type. The outperformance of safe haven assets, itself a manifestation of fear-driven buying, may have an impact on the mispricing of volatility products since it would imply investor excess demand for risk management vehicles. The Bloomberg index for this data series is XAU Curney.

7.4 The cost of insuring against tail risks

Interest in protecting financial interests against tail risks or 'black swan' events is another byproduct of the 2008-9 credit crisis. Note, this is not the same as the previous factor. Investing in safe haven assets is motivated by the return of principal, not necessarily the return on principal, and thus this typically involves a reallocation of capital from one source (higher yielding riskier asset) to another (lower-yielding safer asset). Insuring against tail risks, however, implies continuing to hold the higher-yielding riskier asset, but removing the impact of disaster outcomes via hedging vehicles. The relative cost of insuring against tail risks is a variable that captures this important element of investor attitudes and behavior. My initial conjecture is that rising interest in hedging tail risks exacerbates the forecast bias in the VIX term structure.

As a suitable proxy, I will include the 25-delta 6-month option implied volatility risk reversals for the USDBRL exchange rate. Expressed in units of annualized volatility per annum, this time series tracks the

²⁶The HFRX Global Hedge Fund Index is designed to be representative of the overall composition of the hedge fund universe. It is comprised of all eligible hedge fund strategies; including but not limited to convertible arbitrage, distressed securities, equity hedge, equity market neutral, event driven, macro, merger arbitrage, and relative value arbitrage. The strategies are asset weighted based on the distribution of assets in the hedge fund industry.

²⁷Currency liberalization has helped the Chinese renminbi hold its value despite expectations of a soft/hard landing for the country's economy.

²⁸Hard assets such as forestry, land, and real estate are in general more difficult to procure as investments, manage, and value.

²⁹The Swiss National Bank pegged the franc to the euro on 6-September-2011 in response to excessive inflows and currency strength which it declared to be "a threat to the economy". The level of the peg represented over a 20% depreciation of the franc from recent highs.

differential between out-of-the-money puts and out-of-the-money calls on the Brazilian real (a high-yielding emerging market currency widely considered a risky asset). The greater the differential, the greater the interest in buying protection against a fall in the real. This is associated with rising aversion to risk. The Bloomberg index for this data series is USDBRL25R6M Currency.

7.5 Credit risk

Seeking arbitrage profits generally involves capital intensive strategies³⁰. If the cost of credit rises either too quickly or too far above cost of funds, the speculative capital available for such pursuits will decrease. Swap spreads represent the yield differential, in basis points, for swap contracts versus treasuries. Swap contracts are the benchmark instrument for pricing loans in the private sector, while treasuries represent the US government's cost of borrowing. Divergences between the two are associated with rising levels of credit risk in the financial system. Because the tenor of the trades involving VIX futures are under 1-year, I will include the on-the-run 12-month swap spreads in the regression, available through Bloomberg ticker USSP2 Index.

8. Arbitrage existence and persistence

The premise behind financial markets arbitrage is that asset mispricing or systematic forecast biases may be exploited by simultaneously buying and selling two different versions of the same risk profile, at two different prices. The underpriced (overpriced) asset is bought (sold), and the risk is offset by selling (buying) the other. Barring any unforeseen circumstances³¹, the price differential is the profit. In other words, it is a hedged bet offering lower risk and lower return. I will construct an arbitrage strategy based on the forecast bias documented that is implementable in practice, and evaluate the extent to which changes in the magnitude of the arbitrage profits are impacted by each of the factors discussed. Note the arbitrage profits are a proxy for the forecast bias³². This proxy is necessary for evaluating its existence and persistence in a practical setting.

³⁰This will be illustrated in sections 8.1 and 8.2, which outline the construction of the arbitrage strategy involving the VIX term structure.

³¹Such as financial fraud, insolvency, illiquidity.

³²The terms may be used interchangeably when discussing the results and intuition.

8.1 Replicating the VIX index

The textbook approach for constructing an arbitrage strategy which aims to profit from the forecast bias discussed in this paper would involve systematic selling of VIX futures against a long position in a VIX-replicating portfolio constructed from underlying options on the S&P 500 equity index. As characterized previously, ex ante forecasts of the future level of the VIX extracted from the VIX term structure (constructed from VIX futures prices) consistently overshoot ex post realizations. VIX futures would be sold, and the risk would be hedged by buying the synthetically constructed VIX. The premium earned from the former would be greater than the cost of replicating the index, and the difference would be profit. Although mathematically precise, this avenue is not easily implemented in practice. Even if implementation is attempted, the performance slippage between the VIX index and the VIX-replicating portfolio may become large, thereby eating into the expected return from the arbitrage.

There are three elements that challenge the feasibility of constructing a VIX-replicating portfolio: the volume of trades, the frequency of trade rebalancing required, and the high transaction costs for low-delta options. Strict construction of the definition of the VIX would involve holding portfolios of hundreds, even thousands of options, at any one particular time. Table 1 contains a collection of options on SPX that would be involved in a single hypothetical calculation of the VIX as of 2-Nov-2012. In practice, option strikes are available for every five points on the SPX index, however for expositional purposes, the intervals used are fifteen and twenty-five SPX points apart for near and next-term maturities respectively.

Once the portfolio is constructed, price fluctuations of the SPX index would require rebalancing the VIX-replicating portfolio multiple times a day. The greater the volatility of the SPX index, the more rebalancing required, resulting in thousands of individual transactions a day, often times in just a single hour. Developing a computer algorithm for executing the trades is necessary.

Finally, the theoretical calculation of the VIX uses mid prices. In financial markets, the mid price is the price between the best price of the selling dealer's offer or ask price and the best price of the buying dealer's bid price. Many times it is simply the average or midpoint of the current bid and ask prices being quoted by the dealer. Constructing and managing a VIX-replicating portfolio represents an important deviation from theory, as it is not possible to transact at mid prices. Options would be bought at the dealers offer price (higher than the mid), and options would be sold at the seller's bid price (lower than the mid price). Each trade executed conceivably presents a departure away from the benchmark VIX index. In fact, such transaction costs increase disproportionately for deep out-of-the-money options.

It is possible to deviate from strict construction of the VIX definition outlined by the CBOE. Delisle et

al. (2010) offer an alternative VIX-replicating portfolio which would significantly reduce the administration requirements. However, the tradeoff of such an approach is that it would undoubtedly introduce an additional layer of performance slippage between theory (the VIX index) and practice (the VIX-replicating portfolio).

8.2 Arbitrage strategy, practical approach

In contrast, the results of the hypothesis tests can be used to construct an arbitrage strategy that may be feasibly implemented in practice. Table 3 establishes that the forecast bias is greatest at the short-end of the VIX term structure, and narrowest at the long-end. Tables 4-6 confirm this across sub-periods. The foundation for the arbitrage based on these results involves systematic selling the VIX term structure at the short-end where futures are generally overpriced (leg 1), and hedging this position by establishing a long exposure, equal in notional, to the long-end of the VIX term structure, where the expectations hypothesis is upheld, implying no forecast bias (leg 2). The log return of leg 1 is given by the following

$$r_{leg1} = \sum_{j=1}^2 \ln \left(\frac{\mathcal{F}_{0,j \rightarrow j+1}^{VIX}}{\sigma_{j \rightarrow j+1}^{VIX}} \right), \quad (15)$$

where $\mathcal{F}_{0,1 \rightarrow 2}^{VIX}$ and $\mathcal{F}_{0,2 \rightarrow 3}^{VIX}$ represent the near and next futures contracts that will be shorted at time $t = 0$, the payouts of which will be determined by the prevailing level of the VIX one and two months later, denoted by $\sigma_{1 \rightarrow 2}^{VIX}$ and $\sigma_{2 \rightarrow 3}^{VIX}$ respectively. Leg 1 is expected to generate profits on average.

Similarly, the log return of leg 2 is given by the following

$$r_{leg2} = \sum_{j=1}^7 \ln \left(\frac{\mathcal{F}_{0,j \rightarrow j+1}^{VIX}}{\sigma_{j \rightarrow j+1}^{VIX}} \right), \quad (16)$$

where $\mathcal{F}_{0,j \rightarrow j+1}^{VIX}$ for $j = 1 \dots 7$ represent the long positions in VIX futures established at time $t = 0$, whose payouts will be determined according to the prevailing level of the monthly fixes in the VIX index starting with the near futures expiry and ending at the 7-month expiry date. Leg 2 is also expected to generate profits on average, albeit smaller than leg 1, according to the lower beta coefficients from Table 3.

The profit to the arbitrageur is the net payout from simultaneously buying leg 1 and selling leg 2 at the respective ratios given by

$$r_{arb} = \frac{r_{leg1}}{2} - \frac{r_{leg2}}{7}, \quad (17)$$

where the gains from being long leg 1 are expected to exceed losses from shorting leg 2. The trade is replicated weekly. Changes in the size of the arbitrage profit, which is an expression of the forecast bias identified through the various tests of the expectations hypothesis, will be the dependent variable in the regression.

8.3 Regression results

I explore the contribution of the factors that might influence the size of the forecast bias over time by regressing weekly changes in the realized profit from the arbitrage on weekly changes in the explanatory variables. The regression is carried out twice. The first uses the entire sample of data. The second evaluates only the post-crisis period. Table 7 reports the results.

- Insert Table 7 here -

Starting with the full period analysis, the results indicate that the forecast bias is affected only by the capital flow variable. The sign of the coefficient is particularly illuminating. It suggests that rises in global hedge fund net asset values widen the arbitrage available, itself a representation of the forecast bias. The coefficient estimate indicates that a 10% increase in hedge fund capital in the system would be associated with a 16.9% rise in arbitrage profits. It is important to note that this regressor includes both capital inflows and capital appreciation. This is the definition of net asset value. In other words, it cannot be unequivocally determined that a rise in net asset value directly implies inflows into a particular arbitrage strategy. Rises in net asset value imply both inflows and capital gains.

The sub-period regression indicates that the forecast bias is affected by both the open interest and the capital flow variables, significant at the 5% level. The fit as compared to the full period analysis is superior, based on the improved R-squared reading of 12.1%. The coefficients are both positive, suggesting that capital flows specifically into VIX futures and into the greater market in general via hedge funds, in concert, exacerbate the forecast bias in the VIX term structure. The opposite situation would be expected according to the slow-moving capital hypothesis. Valuable insight is gained from differentiating between the parties involved in these two types of capital allocation. In addition, the interpretation of the capital flow variable must be expanded.

8.3.1 Open interest factor implies inflows from non-professional investors

The transmission mechanism from ETF's to VIX futures markets has been established. In order to provide individual investors with a risk profile that tracks the underlying VIX index, professional money managers offering such ETF's must trade futures and, to a lesser degree, options on the VIX. Mass buying of VIX ETF's then translates into large amounts of capital flowing into futures markets. This is confirmed by the open interest variable. Open interest for VIX futures markets averaged approximately 50,000 contracts, prior to 2009 and prior to the launch of VIX ETF funds. After 2009, there was a sharp rise in open interest. This rise represented multiples above previous levels, no doubt related to the proliferation of VIX ETF offerings following the credit crisis of 2008-9.

With that as background, it can be said that the open interest variable used in the regression is largely a reflection of increases or decreases in VIX ETF volumes. Historically, the bulk of ETF usage has come from non-professional investors, as opposed to hedge funds and institutional asset managers³³. Thus, the open interest factor is a reflection of inflows of non-professional investors, which are in general, less sophisticated than professional investors in a relative sense. This is passive capital in search of buy and hold strategies, as opposed to arbitrage opportunities.

8.3.2 Capital flow factor implies inflows of more sophisticated capital, as well as capital appreciation

The capital flow variable used in the regression is proxied by the HFRX Global Hedge Fund Index, a strategy-weighted representation of the total net asset value of the hedge fund industry comprised of all eligible hedge fund strategies including but not limited to convertible arbitrage, distressed securities, equity hedge, equity market neutral, event driven, macro, merger arbitrage, and relative value arbitrage. This undoubtedly reflects more sophisticated capital which could in theory be readily deployed to exploiting arbitrage opportunities as they arise. A hedge fund manager generally has full discretion of the asset allocation decision. As previously stated, net asset value reflects increases in inflows as well as capital appreciation. Due to this confounder, it can be established that the open interest factor is more direct measure of capital flows into VIX-related strategies.

³³According to the annual study of the United States investment management market conducted by Greenwich Associates during 2010, ETF usage amongst United States pension funds, endowments, and foundations grew to approximately 14%.

8.3.3 Intuition

The positive coefficients for the open interest and the capital flow variables for the regression following the credit crisis of 2008-9 imply that increases in the availability of capital contribute to the biasness of the VIX term structure. The open interest variable reflects inflows of capital from non-professional investors, which are generally less sophisticated, while the capital flow variable reflects inflows of more sophisticated capital, in addition to capital appreciation. Both are happening in tandem. It is the former, the open interest variable, that can be interpreted to be evidence against the slow-moving capital hypothesis, as it more clearly reflects inflows into VIX-related strategies. This phenomenon is illustrated in greater detail in Figure 8.

- Insert Figure 8 here -

I construct a series based on the 2-month VIX term structure tests of the expectations hypothesis. The points on the time series represent the beta coefficients for rolling 1-year regressions, per equation (14). A decrease in the coefficient away from one implies a larger forecast bias. The series is then plotted against the market capitalization for the VXX ETF, the largest offering from this family of funds. The link between the size of the forecast bias and the size of the VIX fund is evident, as previously established from tests of the expectations hypothesis.

Lastly, the positive coefficient for the capital flow variable may also imply that as hedge funds perform well, so does the market in general. The combination of a bullish cycle for risk assets and more inflows into VIX futures (as suggested by the open interest variable) presents the perfect storm for a widening of the forecast bias in the VIX term structure. As the S&P 500 rises, the VIX index falls. If this is happening at the same time as strong inflows are heading into the market for VIX futures, this will result in a steepening of the VIX term structure, as shown, and translate into higher forecasts for the future level of the VIX.

9. Closing comments

The focus of this paper is first to evaluate the information content of VIX futures prices, the core method of attaining exposure to the VIX index. There are two main objectives. The first is to identify, characterize, and quantify the forecast bias of the VIX term structure, which is constructed from VIX futures prices. There are four key findings along these lines.

First, VIX futures are consistently overpriced relative to the subsequent moves in the underlying VIX index. Second, the forecast bias increases substantially with shorter tenors. Third, the forecast bias is smaller comparatively during periods of extreme volatility. Four, deviations from expectations hypothesis are greatest during the years following the US credit crisis of 2008-9. These findings describe the VIX-VIX Futures Puzzle.

The second objective, once the forecast bias has been documented, is to shed light on why the forecast bias exists and why it persists. I first identify a number of factors that might influence the size of the forecast bias over time: futures open interest, hedge fund capital flows, performance of safe haven assets, the costs of insuring against tail risks, and the amount of credit risk in the financial system. I then construct an arbitrage strategy which aims to profit from the forecast bias I identify, and regress weekly realized changes in the arbitrage profits versus changes in the factors. Note the size of the arbitrage profits are a proxy for the size of the forecast bias. The results of this regression for the sub-period following the 2008-9 credit crisis suggest that capital inflows into VIX futures (the open interest factor) and into the greater market in general via hedge funds (the capital flows factor), in concert, exacerbate the forecast bias in the VIX term structure. This is a particularly surprising result, as the opposite situation would be suggested by application of the slow-moving capital hypothesis. The persistence of the arbitrage is related to key differences between the parties involved in the capital allocations. The open interest variable reflects inflows from non-professional investors, which are generally less sophisticated than professional investors (and thus not typically looking for arbitrage profits), while the capital flow variable reflects inflows of more sophisticated capital, as well as capital appreciation. As the VIX has become more accessible to the average investor, this has created distortions to VIX futures markets, especially for short-dated tenors which are most actively used for management of ETF funds. Lastly, the positive coefficient for the capital flow variable may also imply that as hedge funds perform well, so does the market in general. The combination of a bullish cycle for risk assets and more inflows into VIX futures (as suggested by the open interest variable) presents the perfect storm for a widening of the forecast bias in the VIX term structure.

Future research will be dedicated to evaluating alternative explanations of the VIX-VIX Futures puzzle, including the development of the idea that the forecast bias in the VIX term structure may be attributed to investor appetite for paying a certain premium to VIX products because of the diversification benefits offered. Another interesting avenue for research involves the development of a new formulation for testing EH for the VIX and other volatility products, deviating from a risk-neutral framework. In addition, while I focused on tests of the expectations hypothesis applicable to the term structure of implied volatility, it

should be noted that VIX futures are more of a hybrid product. By construction, the VIX is determined from underlying option prices, however, VIX futures themselves offer linear returns, a feature of futures and forward contracts. Methods for testing the expectations hypothesis in futures markets were not directly applied here since the VIX is not an investable asset such as gold, oil, and other commodities where cost of carry relationships, storage costs, transportation costs, and insurance premiums impact futures pricing. Nonetheless a new formulation combining the elements of both futures and options is worth exploring.

Appendix

A. VIX portfolios: Hedge or diversification?

The objective is to establish that capital allocations to VIX portfolios provide diversification, not direct offsets aimed at covering losses, a feature typically associated with hedging activity. This class of investor desires a risk profile that is directly tied to market sentiment and not necessarily to underlying equity prices. There is a common misconception that a capital allocation to a portfolio constructed from VIX products would provide a suitable hedge to a long equity position given the strength of the co-movement between the VIX and the SPX on a mark-to-market basis. Figure 9, however, shows a strong counterargument.

- Insert Figure 9 here -

A dollar invested in the S&P 500 index in February 2011 would have been worth about the same in February 2012, with minimal divergence from the benchmark starting point. A dollar invested in VXX, an exchange traded fund (ETF) which takes long positions in VIX futures, would also have been worth a dollar at the end of the period but its value would have, in contrast, oscillated significantly over this time period, falling over thirty percent in the summer of 2011 and then nearly doubling in the fourth quarter of 2011 when the Euro zone sovereign debt crisis dominated headlines. The combined risk profile of these two assets would not constitute prudent hedging activity in a traditional sense. A hedge by definition should offset losses, not enhance returns or exacerbate volatility as is evident in this case.

Figure 10 shows the rolling three-month correlation of changes between the SPX and the VIX, and the realized standard deviation of each.

- Insert Figure 10 here -

As established, the two indices are well correlated. Correlation however does not tell the whole story, and a high correlation, even if persistent, does not ensure that gains and losses on the core investment will be sufficiently offset by gains and losses on the hedge. The optimal hedge ratio may be changing despite the stability of the correlation.

The time-varying nature of optimal hedge ratios has been addressed extensively in the literature. There is no debate about the fact that optimal hedge ratios change across time, the controversy centers around whether or not there is merit in implementing dynamic versus static approaches in the context of portfolio construction. Figlewski (1984), Lypny (1988), and Baillie and Myers (1989) find that a dynamic hedge

strategy outperforms the time invariant hedge, while Smirlock (1985) and Cecchetti et al. (1988) find that time invariant hedge ratios perform better than the dynamic approach.

I will sketch out the basic theoretical argument for explaining the ineffectiveness of the VIX as a true hedge for a long position in the SPX. Despite the strong correlation, the ratio of the volatilities is the source of the ineffectiveness of the hedge. Consider a two-asset portfolio comprised of long positions in the SPX and VIX indices³⁴ whose combined variance is given by

$$\sigma_{SPX,VIX}^2 = w_{SPX}^2 \sigma_{SPX}^2 + w_{VIX}^2 \sigma_{VIX}^2 + 2\rho w_{SPX} w_{VIX} \sigma_{SPX} \sigma_{VIX}, \quad (18)$$

where w represents the weights of each asset, σ^2 are the respective variances, and ρ is the correlation between the two assets. The two assets by definition offer linear returns. Suppose w_{SPX} is normalized at 1.0. The exercise involves selecting w_{VIX} such that the combined portfolio variance reduction is maximized, and thus we compute the partial derivative of equation (18)

$$\frac{\partial \sigma_{SPX,VIX}^2}{\partial w_{VIX}} = 2w_{VIX} \sigma_{VIX}^2 + 2\rho \sigma_{SPX} \sigma_{VIX}. \quad (19)$$

The optimal allocation to the VIX is then found by setting (20) equal to zero and solving for w_{VIX}

$$w_{VIX} = -\rho \frac{\sigma_{SPX}}{\sigma_{VIX}}. \quad (20)$$

Equation (21) says that the optimal allocation to the VIX should be proportional to the ratio of the volatilities of both elements in the portfolio. Figure II.11 shows this ratio historically, based on the rolling 3-month calculations displayed in figure 11.

- Insert Figure 11 here -

Working with the assumption that $\rho = 1$, a reading of 0.10 as in 2006 would suggest the optimal allocation to the VIX should be 10 cents for every dollar invested in the SPX, in other words, equal to the ratio depicted in figure 6 since w_{SPX} is normalized at 1.0. During the credit crisis of 2008, this optimal hedge ratio would have risen close to five times 2006 levels. The intuition for this is as follows. By construction, the VIX is designed to maintain a constant price sensitivity to the implied volatility of the SPX over a one-month period. This is achieved by constructing the index as being a weighted average of all OTM options for the

³⁴For expositional purposes, I will assume the VIX is directly investable. The results should hold if actual market instruments are used.

two nearby expiries, centered around the prevailing SPX futures curve, at any point in time³⁵. Thus, changes in the value of the underlying SPX index will change the hypothetical portfolio of options that determine the value of the VIX. As the SPX drops, the hedger would ideally want to continue to hold options struck at the higher strikes. The VIX, however, does not offer this risk profile, as it is only defined by the full set of OTM options. Once an option goes in-the-money (ITM), it will no longer be reflected in the price of the VIX. Instead, the VIX holder would synthetically own a greater concentration of the lower strikes in tandem with the fall in the SPX index, and none of the higher ITM strikes that would insulate the fall in equity prices. Maintaining an optimal hedge ratio would then require upsizing w_{VIX} as spot SPX is dropping. Along similar lines, as SPX rises, maintaining an optimal hedge ratio would require reducing w_{VIX} . This is seen clearly confirmed by figure 11, the optimal volatility-reducing holding of the VIX, w_{VIX} , would have fallen dramatically following the post 2008-9 credit-crisis bottom of the SPX in March 2009.

Furthermore we can establish that a dramatic shift in the ratio of volatilities has the same impact as that of a significant drop in correlation from the standpoint of hedge effectiveness, defined as

$$\mathcal{H}_{VIX} = -\frac{\sigma_{SPX,VIX}^2 - \sigma_{SPX}^2}{\sigma_{SPX}^2}, \quad (21)$$

where \mathcal{H} takes on a maximum value of 1.0 denoting perfect effectiveness. Combining arrives at the expression for hedge effectiveness

$$\mathcal{H}_{VIX} = \rho. \quad (22)$$

In other words, if the hedge ratio is constructed according to the optimal ratios outlined, then \mathcal{H} is purely a function of the correlation. Suppose however that due to a change in the ratio of the volatilities as depicted in figure 11, the portfolio is hedged according to the less optimal hedge ratio given by

$$w_{VIX} = -\frac{1}{3}\rho\frac{\sigma_{SPX}}{\sigma_{VIX}}, \quad (23)$$

which says that despite the rise in ratio of volatilities, the allocation to the VIX, w_{VIX} , was not changed. The hedge effectiveness for this less optimal portfolio would be given by

$$\mathcal{H} = \frac{1}{18}\rho^2. \quad (24)$$

³⁵Refer to table 1 for illustration.

Figure 12 contains a plot of the efficiency under the optimal and sub-optimal hedge ratios against all possible values of ρ .

- Insert Figure 12 here -

It is clear that even at perfect correlation levels, the hedge efficiency is low under a sub-optimal hedge ratio. Otherwise stated, the impact on hedge efficiency of a portfolio of two assets that exhibit a high correlation to one another but constructed at a suboptimal hedge ratio is similar to that of a portfolio of assets that exhibit a low correlation, even if this portfolio is constructed according to the optimal hedge ratio. The rapidly changing ratio of the volatilities is the root of hedge inefficiency. Low hedge efficiency is associated with diversification, not hedging.

A true hedge from the standpoint of offsetting gains and losses on an underlying core investment in S&P 500 stocks would involve SPX futures or options. Futures and options span the full spectrum of payoff outcomes. The former offers linear returns, full downside protection, does not require upfront premium, but offers no upside potential. The resulting risk profile for hedging via SPX futures is equivalent to exiting the investment in the SPX altogether. On the other hand, options on the SPX offer asymmetric returns, full protection and full upside participation, but require initial premium which averages five to ten percent per annum. VIX portfolios are generally marketed as combining the attractive elements of both SPX futures and options: linear returns, no upfront premium and upside potential. There is, however, no free lunch in financial markets and thus the holder must give up having full downside protection. In addition, the value of the VIX portfolio itself may fall in tandem with rises in equity prices. Investor incentives with regard to allocations to VIX portfolios are thus associated with diversification.

B. Tests of expectations hypothesis for VIX

The VIX, widely considered the market's fear gauge, is an index which represents the market's estimate of the realized volatility in the S&P 500 index (SPX) over a one month period. Despite its literal definition, the market's view of future realized volatility over the short-term, the payoffs of VIX products is not tied to the manifestation of realized volatility over this future period. Payoffs on the aforementioned VIX-linked products are instead tied to futures changes in the market's estimate of future volatility. This is an important point which is the source of some confusion as it relates to use and interpretation.

This test evaluates the accuracy of the VIX as an ex ante estimate of future realized volatility in the SPX. Figure 13 plots the VIX at time t versus realized volatility in the SPX at $t + 1$. Instances where the VIX is higher (lower) imply situations where the ex ante forecast was greater (lower) than the ex post realization.

- Insert Figure 13 here -

Table 8 contains the results of a test I carried out using the data in this paper which assesses the forecast accuracy of the VIX. Note the main paper addresses VIX futures, this is a separate question altogether.

- Insert Table 8 here -

In general, ex ante estimates overshoot ex post realized volatility in the S&P 500. Beta coefficients below one imply VIX ex ante estimates overshoot ex post realizations of S&P 500 realized volatility. The highest coefficients were recorded in 2009, a period of great financial distress. In general, these results are consistent with Fleming (1998), Bakshi and Kapadia (2002), Buraschi and Jackwerth (2001), Coval and Shumway (2001), and Pan (2002).

References

- [1] Ait-Sahalia, Y., Karaman, M., Mancini, L. (2012). The term structure of variance swaps, risk premia and the expectations hypothesis. Working paper. Princeton University, NBER, Swiss Finance Institute, University of Zurich, and EPFL.
- [2] Bakshi, G., Kapadia, N. (2002). Delta hedged gains and the negative market volatility risk premium. *Review of Financial Studies* 16, 2, 527-566.
- [3] Baillie, R., Myers, R. (1989). Modeling commodity price distributions and estimating the optimal futures hedge. Unpublished manuscript. Department of Economics, Michigan State University.
- [4] Becker, R., & Clements, A. (2008). Are combination forecasts of S&P 500 volatility statistically superior? *International Journal of Forecasting*, 24, 122- 133.
- [5] Becker, R., Clements, A. & White, S. (2006). Does implied volatility provide any information beyond that captured in model-based volatility forecasts? *Journal of Banking and Finance*.
- [6] Benzoni, L. (2001). Pricing options under stochastic volatility: An empirical investigation. Working paper, Carlson School of Management, Minneapolis.
- [7] Blair, B., Poon, S. H., and Taylor, S. (2001). Forecasting S&P 100 volatility: The incremental information content of implied volatilities and high frequency index returns. *Journal of Econometrics* 105, pp. 5-26.
- [8] Bollerslev, T., Tauchen, G., Zhou, H. (2009). Expected stock returns and variance risk premia. *Review of Financial Studies* 22, 4463.
- [9] Breeden, D., Litzenberger, R. (1978). Prices of state-contingent claims implicit in options prices. *Journal of Business* 51, 621-651.
- [10] Britten-Jones, M., Neuberger, A. (2000). Option prices, implied price processes, and stochastic volatility. *Journal of Finance* 55, 839-866.
- [11] Brunnermeier, M., Pedersen, L. (2009). Market liquidity and funding liquidity. *Review of Financial Studies* 22, 2201-2238.
- [12] Buraschi, A., Jackwerth, J. (2001). The price of a smile: Hedging and spanning in option markets. *Review of Financial Studies* 14, 2, 495-527.
- [13] Byoun, S., Kwok, C., and Park, H. (2003). Expectations hypothesis of the term structure of implied volatility: Evidence from foreign currency and stock index options. *Journal of Financial Econometrics* 1, 126-151.
- [14] Campa, J., Chang, K. (1995). Testing the expectations hypothesis on the term structure of volatilities in foreign exchange options. *Journal of Finance* 50, No. 2.
- [15] Campbell, J.Y., Lo, A.W. & MacKinlay, A.G. (1997). *The Econometrics of Financial Markets*. Princeton University Press. Princeton NJ.
- [16] Carr, P., Lee, R. (2009). Volatility derivatives. *Annual Review of Financial Economics* 1, 319-339.
- [17] Carr, P., Wu, L. (2006). A tale of two indices. *Journal of Derivatives* 13, 13-29.
- [18] Cecchetti, S., Cumby, R., Figlewski, S. (1988). Estimation of the optimal futures hedge. *Review of Economics and Statistics*, 70, 623-630.
- [19] Chernov, M. (2001). Implied volatilities as forecasts of future volatility, the market risk premia, and returns variability. Working paper, Columbia Business School.

- [20] Coval, J., Shumway, T. (2001). Expected option returns. *Journal of Finance* 56, 3, 983-1009.
- [21] Daigler, R., Rossi, L. (2006). A portfolio of stocks and volatility. *Journal of Investing*. Summer, 99-106.
- [22] Delisle, J., Doran, J., Krieger, K. (2010). Volatility as an asset class: Holding VIX in a portfolio. Unpublished manuscript.
- [23] Demeterfi, K., Derman, E., and Kamal, M. (1999). A guide to volatility and variance swaps. *Journal of Derivatives* 6, 9-32.
- [24] Dreschsler, I., Yaron, A. (2011). What's vol got to do with it? *Review of Financial Studies* 24, 1.
- [25] Dupire, B. (2006). Model free results on volatility derivatives. Working paper. Bloomberg, NY: SAMSI.
- [26] Fama, E. (1984). The information in the term structure. *Journal of Financial Economics* 13, 509-528.
- [27] Figlewski, S. (1986). Hedging with financial futures for institutional investors. Ballinger Publishing: Cambridge, MA.
- [28] Fleming, J. (1998). The quality of market volatility forecasts implied by S&P 100 index option prices. *Journal of Empirical Finance* 5, 317-345.
- [29] Froot, K. (1989). New hope for the Expectations Hypothesis of the term structure of interest rates. *Journal of Finance*. Vol. XLIV, No. 2.
- [30] Gouriéroux C. & Jasiak J. (2001). *Financial Econometrics*. Princeton University Press. Princeton NJ.
- [31] Gromb, D., Vayanos, D. (2002). Equilibrium and welfare in markets with financially constrained arbitrageurs. *Journal of Financial Economics* 66, 361-407.
- [32] Hull, J., White, A. (1987). The pricing of options on assets with stochastic volatility. *Journal of Finance* 42, 281-300.
- [33] Jiang, G., Tian, Y. (2007). Extracting model-free volatility from option prices: An examination of the VIX index. *Journal of Derivatives* 14(3), 1-26.
- [34] Konstantinidi, E., Skiadopoulos, G. (2011). Are VIX futures prices predictable? An empirical investigation. *International Journal of Forecasting*, forthcoming.
- [35] Konstantinidi, E., Skiadopoulos, G., Tzagkaraki, E. (2008). Can the evolution of implied volatility be forecasted? Evidence from European and US implied volatility indices. *Journal of Banking and Finance* 32, 2401-2411.
- [36] Lin, Y. (2007). Pricing VIX futures: Evidence from integrated physical and risk-neutral probability measures. *Journal of Futures Markets* 27, 1175-1217.
- [37] Lu, Z., Zhu, Y. (2010). Volatility components: The term structure dynamics of VIX futures. *Journal of Futures Markets* 30, No. 3, 230-256.
- [38] Luo, X., Zhang, J. (2012). The term structure of VIX. *Journal of Futures Markets*, 32(12), 1092-1123.
- [39] Lypny, G. (1988). Hedging foreign exchange risk with currency futures: Portfolio effects. *Journal of Futures Markets* 8, 703-715.
- [40] Matthias F., Longstaff, F., and Lustig H. (2013). The TIPS-Treasury Bond Puzzle. DOI: 10.1111/jofi.12032.
- [41] Mixon, S. (2007). The implied volatility term structure of stock index options. *Journal of Empirical Finance* 14, No. 3. 333-354.

- [42] Newey, W., West, K. (1987). Hypothesis testing with efficient method of moments estimation. *International Economic Review* 28, 777-787.
- [43] Nossman, M., Wilhelmsson, A. (2008). Is the VIX futures market able to predict the VIX Index? A test of the expectations hypothesis. *Journal of Alternative Investments*.
- [44] Pan, J. (2002). The jump risk premia implicit in options: Evidence from an integrated time-series study. *Journal of Financial Economics* 63, 1, 3-50.
- [45] Poon, S., Granger, C. (2003). Forecasting volatility in financial markets: A review. *Journal of Economic Literature*. Vol XLI, 478-539.
- [46] Poteshman, A. (2001). Underreaction, overreaction, and increasing misreaction to information in the options market. *Journal of Finance* 56, 851-876.
- [47] Sepp, A. (2008b). Pricing options on realized variance in Heston model with jumps in returns and volatility. *Journal of Computational Finance* 11(4), 33-70.
- [48] Shiller, R. (1979). The volatility of long-term interest rates and expectations models of the term structure. *Journal of Political Economy*, Vol. 87, No. 6, 1190-1219.
- [49] Shu, J., Zhang, J. (2012). Causality in the VIX futures markets. *Journal of Futures Markets* 32, No. 1, 24-46.
- [50] Smirlock, M. (1985). An analysis of cross-hedging with Treasury bill futures: Bank specific evidence. Unpublished manuscript, Federal Reserve Bank of Philadelphia.
- [51] Stein, J. (1989). Overreactions in the options markets. *Journal of Finance* 44, 1011-1023.
- [52] Szado, E. (2009). VIX futures and options - A case study of portfolio diversification during the 2008 financial crisis. Working paper, University of Massachusetts, Amherst.
- [53] Taylor, S. J. (2005). *Asset Price Dynamics, Volatility, and Prediction*. Princeton University Press. Princeton NJ.
- [54] Zhang, J., Huang, Y. (2010). The CBOE S&P 500 three-month variance futures. *Journal of Futures Markets* 30(1), 48-70.
- [55] Zhang, J., Zhu, Y. (2006). VIX futures. *Journal of Futures Markets* 26, 521-531.
- [56] Zhu, Y., Zhang, J. (2007). Variance term structure and VIX futures pricing. *International Journal of Theoretical and Applied Finance* 10(1), 111-127.

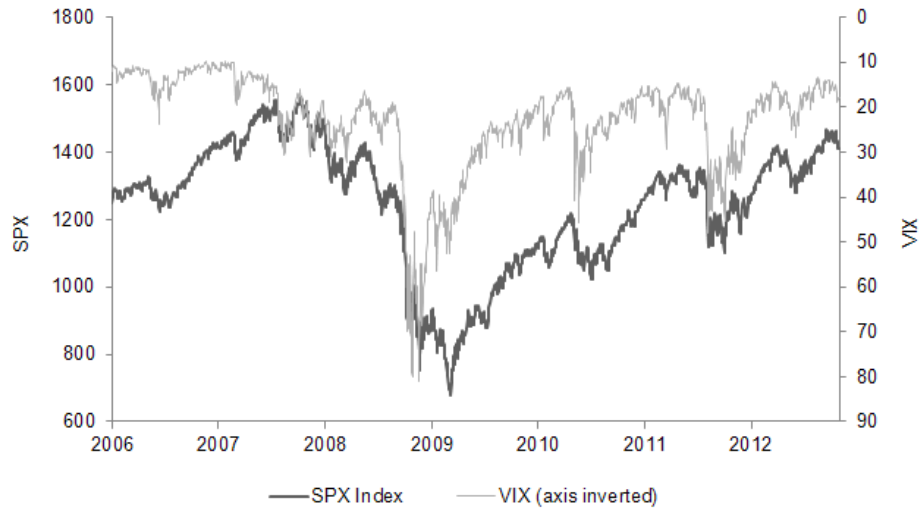


Figure 1. Historical prices for VIX and SPX. This figure tracks historical price data on the VIX and the SPX, from January 2006 to October 2012. The strong association between price series is supported quantitatively by a correlation coefficient, based on daily changes, of -76% for the entire period. In addition, the average three-month rolling correlation is -83%. The negative number implies co-movements occur in the opposite direction.

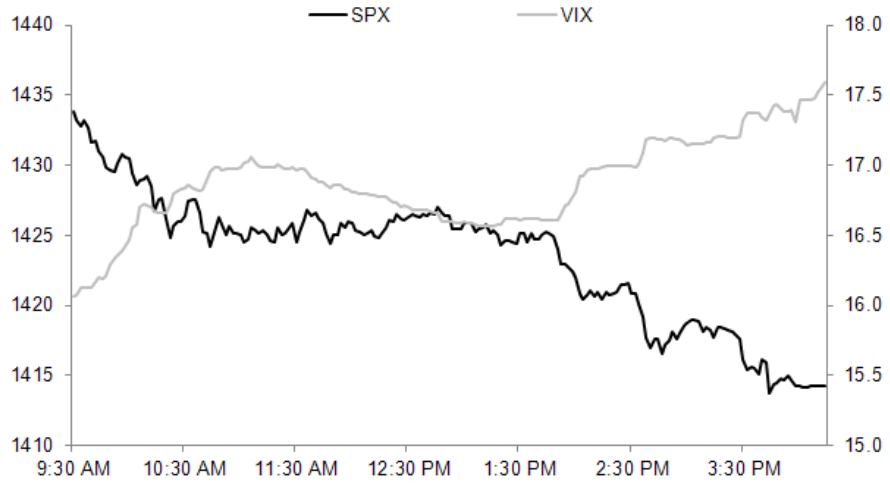


Figure 2. Intra-day prices on VIX and SPX on 2-Nov-2012. This figure tracks historical price data on the VIX and the SPX on an intra-day basis. Although not perfect, the strength of the co-movement is evident.




Figure 3. VIX, VIX futures, and expectations of SPX realized volatility illustrated. This figure illustrates the interaction between the VIX, VIX futures, and expectations of SPX realized volatility. On 1-February-2012, time t , the closing price quotes, in annualized standard deviation terms, for the VIX and the 1st and 2nd futures were $\sigma_t^{VIX} = 18.55$, $\mathcal{F}_{t,j \rightarrow j+1}^{VIX} = 19.85$, and $\mathcal{F}_{t,j+1 \rightarrow j+2}^{VIX} = 22.05$ respectively. Note that σ_t^{VIX} is a forecast of realized volatility of the SPX over the immediate future period, while $\mathcal{F}_{t,j \rightarrow j+1}^{VIX}$ and $\mathcal{F}_{t,j+1 \rightarrow j+2}^{VIX}$ are forecasts of the future forecasts of realized volatility of the SPX. The horizontal dashed arrows represent the period over which such expectations apply.

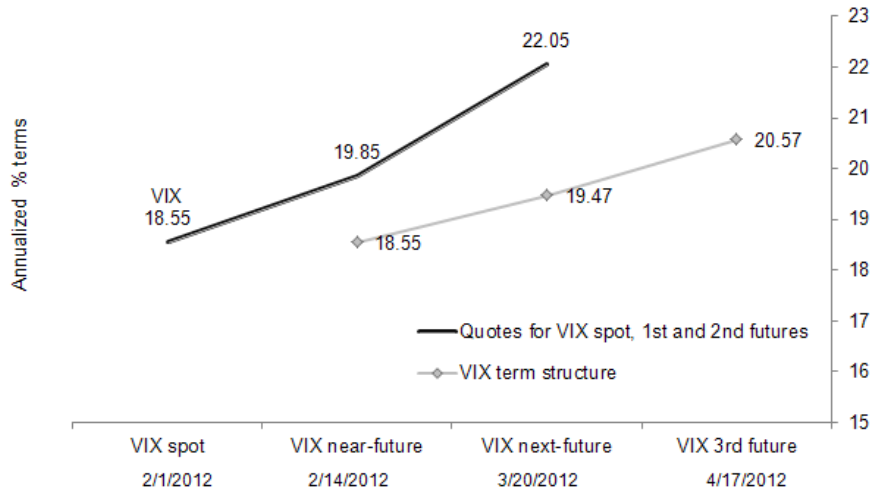


Figure 4. VIX, VIX futures, and resulting VIX term structure. Figure shows the resulting VIX term structure based on the sample quotes from an earlier section. The constructed VIX term structure will be lower than the VIX futures curve when the slope of the VIX futures curve is positive. Forecasts of cumulative volatility become more accurate as the period of time over which volatility is generated grows due to cancellation of error and mean reversion dynamics. Thus a six-month long-dated VIX contract or variance swap represents the expected variance in the SPX over a six-month period, while the six-month VIX future represents the estimate today for the VIX index, or equivalently the thirty-day estimate of the variance in the SPX, six months from now. The latter naturally carries greater uncertainty, and thus a higher volatility price.

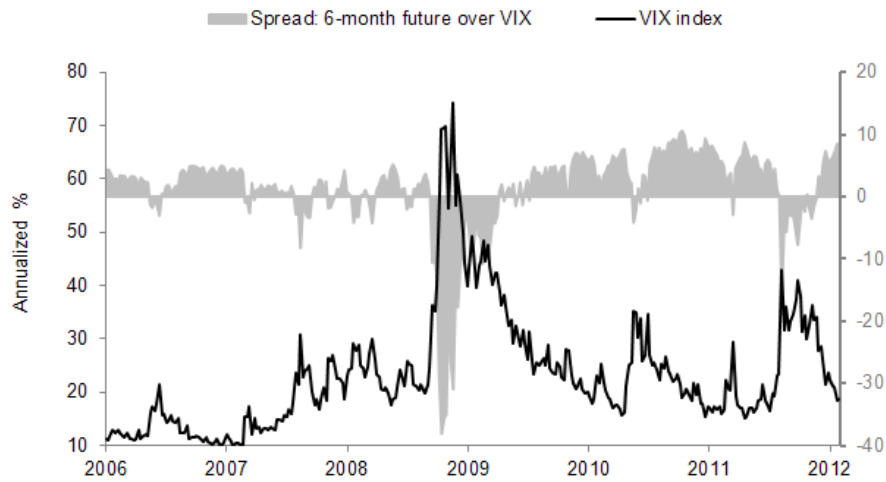


Figure 5. Spread between six month VIX futures versus the VIX. Figure tracks the spread between six-month VIX futures (left axis) and the underlying VIX (right axis) for the entire sample period. A few observations worth noting. For one, roughly 75% of the time, the spread between the long-end future and the VIX is positive. Also, the spread went steeply negative during the height of the 2008 credit crisis and then materially negative in the latter half of 2011 when Euro zone crisis fears peaked. And three, the average spread when positive was substantially greater post-crisis, as opposed to pre-crisis.

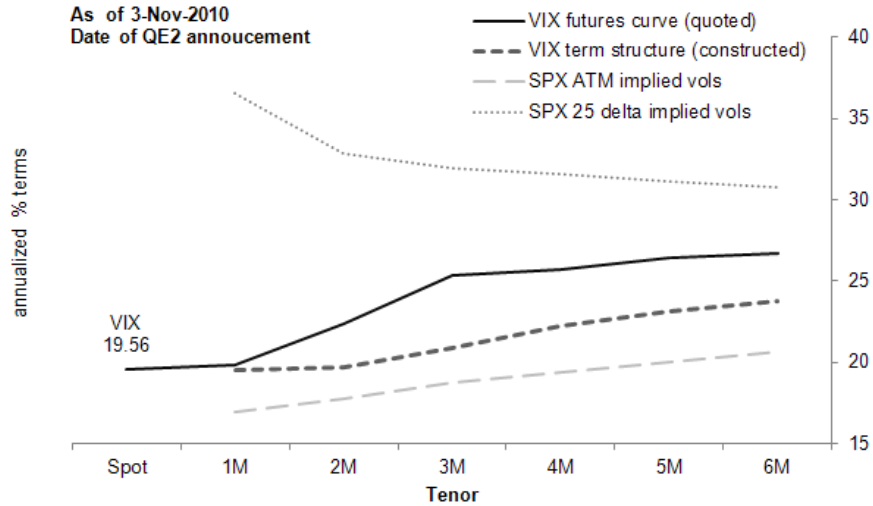


Figure 6. VIX futures curve, VIX term structure, and SPX implied volatility term structure.

Figure is a snapshot of the various curves associated with the VIX as of the close of business on 3-November-2010, the day the second round of quantitative easing (QE2) was announced in the United States. The VIX index closed at 19.56, while the 6-month VIX future closed above 25, displaying a significant premium. As discussed, when the VIX futures curve is upward sloping, the constructed VIX term structure or variance swap curve, will generally be lower. Also depicted in the figure are the term structures for ATM and 25-delta OTM implied volatilities for SPX options. SPX options are the building blocks of the VIX index. It is common for the term structure of ATM volatilities on underlying SPX options to be lower than the VIX term structure across tenors, as the latter prices in the skew that is prevalent for OTM options on all financial assets. By construction, the VIX gives the holder exposure to the full set of OTM options on SPX at any one particular time, and thus the skew is reflected in the higher term structure. Regarding the relationship between the VIX term structure and the term structure of 25-delta OTM SPX options, we would expect the latter to be higher on average. The payoff of a variance swap such as the VIX is convex in volatility. This means that an investor who is long a variance swap will benefit from boosted gains and discounted losses, a phenomenon which is amplified when volatility skew is steep. Thus, the fair strike of a variance swap is often in line with the implied volatility of 40-delta SPX puts, which is lower than that of 25-delta puts in the presence of skew.

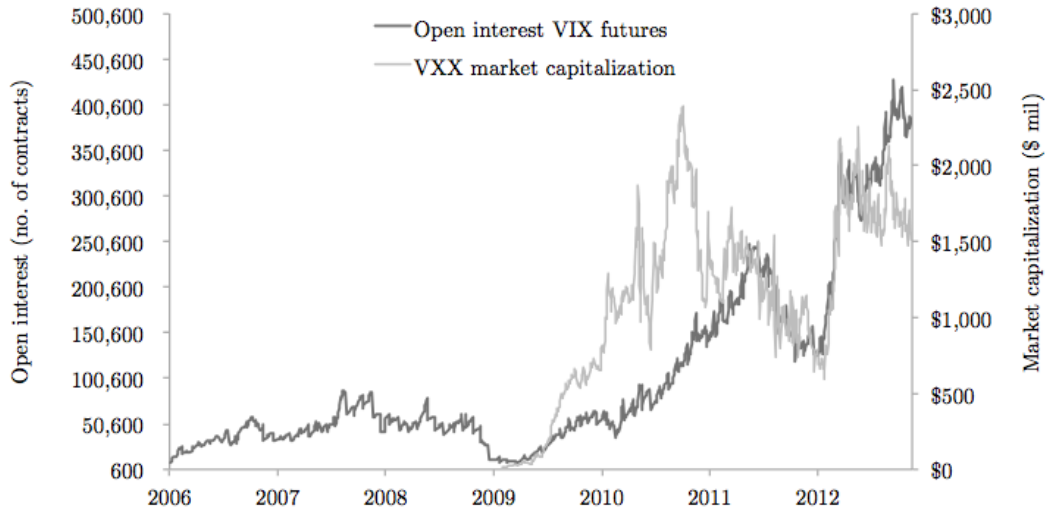


Figure 7. Open interest in VIX futures versus market capitalization on VXX ETF. Figure tracks the open interest in VIX futures along with the market capitalization of the VXX ETF. Open interest averaged approximately 50,000 contracts prior to 2009. After 2009, there was a sharp rise in open interest, no doubt related to the proliferation of VIX ETF offerings. Today, ETF's represent the main vehicle for accessing the VIX, and inflows into ETF's implies inflows into the market for VIX futures.

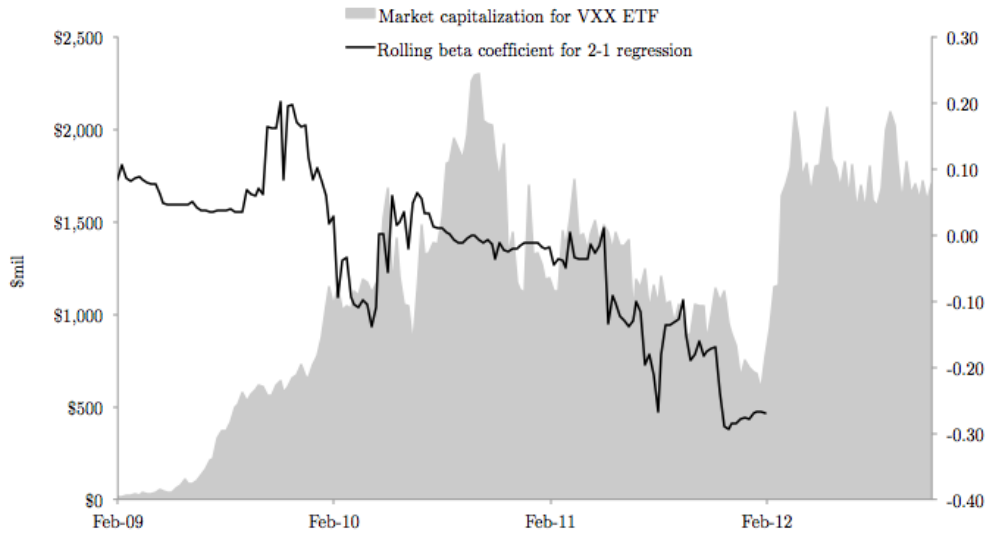


Figure 8. Rolling beta for 2-1 regression versus VXX ETF fund market capitalization. Figure shows that the forecast bias in the VIX term structure has increased following the 2008-9 credit crisis. The points on the time series represent the beta coefficients for rolling 1-year regressions, per equation (14). A decrease in the coefficient away from one implies a larger forecast bias. The series is then plotted against the market capitalization for the VXX ETF, the largest offering. The link between the size of the forecast bias and the size of the VIX fund is evident.

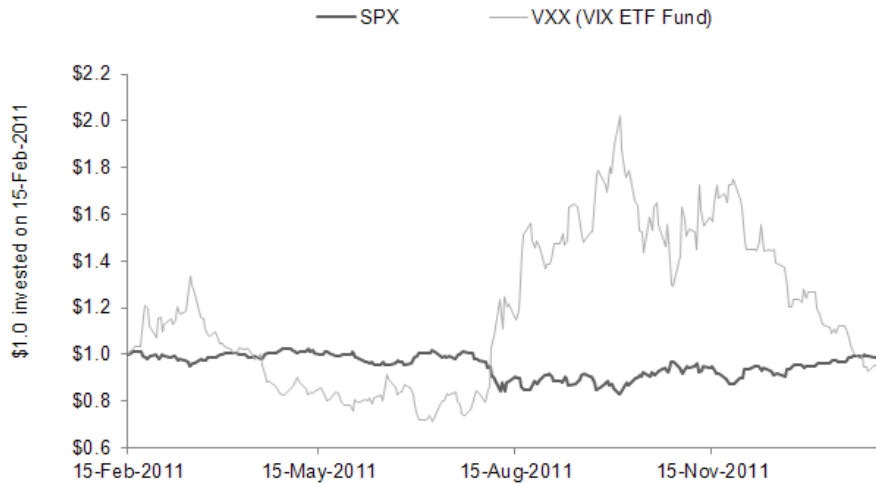


Figure 9. Performance of a dollar invested in the SPX index versus a dollar invested in a VIX ETF. Figure shows that a dollar invested in the S&P 500 index in February 2011 would have been worth about the same in February 2012, with minimal divergence from the benchmark starting point. A dollar invested in VXX, an exchange traded fund (ETF) which takes long positions in VIX futures, would also have been worth a dollar at the end of the period but its value would have, in contrast, oscillated significantly over this time period, falling over thirty percent in the summer of 2011 and then nearly doubling in the fourth quarter of 2011 when the Euro zone sovereign debt crisis dominated headlines. The combined risk profile of these two assets would not constitute prudent hedging activity in a traditional sense. A hedge by definition should offset losses, not enhance returns or exacerbate volatility as is evident in this case.

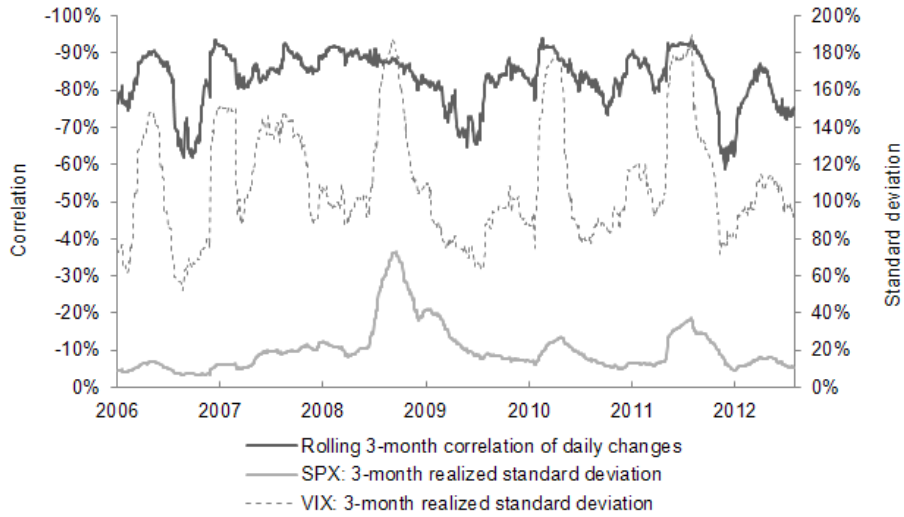


Figure 10. Rolling 3-month correlation and realized standard deviation for SPX and VIX. Figure shows the rolling three-month correlation of changes between the SPX and the VIX, and the realized standard deviation of each. As established, the two indices are well correlated.




Figure 11. Ratio of 3-month realized volatility of SPX over the realized volatility of the VIX. In a two-asset portfolio made of up an underlying investment in the SPX and a hedging vehicle such as the VIX, the optimal allocation to the VIX should be proportional to the ratio of the volatilities of both elements in the portfolio. Figure shows this ratio historically, based on the rolling 3-month calculations of the volatilities.

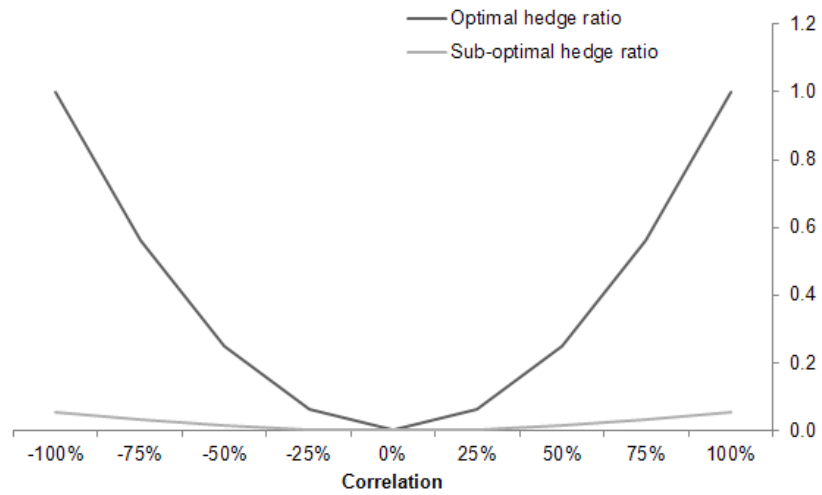


Figure 12. Hedge efficiency at different levels of correlation. Figure contains a plot of the efficiency under the optimal and sub-optimal hedge ratios against all possible values of the correlation coefficient. It is clear that even at perfect correlation levels, the hedge efficiency is low under a sub-optimal hedge ratio. Otherwise stated, the impact on hedge efficiency of a portfolio of two assets that exhibit a high correlation to one another but constructed at a suboptimal hedge ratio is similar to that of a portfolio of assets that exhibit a low correlation, even if this portfolio is constructed according to the optimal hedge ratio.




Figure 13. VIX at time t versus realized volatility in S&P 500 at $t + 1$. Figure plots the VIX at time t versus realized volatility in the SPX at $t + 1$. Instances where the VIX is higher (lower) imply situations where the ex ante forecast was greater (lower) than the ex post realization.

Table 1

Collection of SPX options involved in σ_{VIX} calculation*

Table contains a collection of options on SPX that would be involved in a hypothetical calculation of the VIX as of 2-Nov-2012. In practice, option strikes are available for every five points on the SPX index, however for expositional purposes, the intervals used are fifteen and twenty-five SPX points apart for near and next-term maturities respectively. The salient information to be extracted from the table is as follows: 1) The strike of 1415 is the strike where the price difference between calls and puts for both maturities is smallest, 2) the strike of 1415 will be used to determine F and K_0 , which in turn determines the set of K_i , 3) the range of option strikes used for each maturity will vary as the calculation leaves out options for which the bid price is zero, and 4) the exact collection of options used will change in tandem with changes in the underlying price of the SPX as in-the-money (ITM) options are left out of the calculation.

Date: 2-Nov-2012			SPX range: 1414-1433			σ^{VIX} range: 16.10-17.59		
Near term expiry 17-Nov-2012			Near term expiry 19-Dec-2012					
$F_1 = 1414.5$			$F_2 = 1412.25$					
Strike	Call / Put	Mid price	Strike	Call / Put	Mid price			
1235	P	0.20	1115	P	0.50			
1250	P	0.40	1140	P	0.85			
1265	P	0.60	1165	P	1.40			
1280	P	0.80	1190	P	1.85			
1295	P	1.00	1215	P	2.25			
1325	P	1.25	1240	P	3.30			
1310	P	1.50	1265	P	4.50			
1340	P	2.25	1290	P	6.00			
1355	P	3.45	1315	P	8.50			
1370	P	5.00	1340	P	12.0			
1385	P	8.00	1365	P	16.5			
1400	P	12.0	1390	P	23.3			
1415	P	17.5	1415	P	32.8			
1415	C	17.0	1415	C	31.0			
1430	C	11.0	1440	C	17.0			
1445	C	6.00	1465	C	8.50			
1460	C	3.00	1490	C	4.00			
1475	C	1.50	1515	C	1.80			
1490	C	0.80	1540	C	0.80			
1505	C	0.50	1565	C	0.60			
1520	C	0.40	1590	C	0.40			
1535	C	0.30	1615	C	0.30			
1550	C	0.20	1640	C	0.10			

* Calculation leaves out options for which the bid price is zero

Table 2

Descriptive statistics for VIX, VIX futures, and the resulting term structure (2006-2012)

Table contains descriptive statistics on the VIX, VIX futures, and the resulting VIX term structure spanning the period 2006-2012. The average figures suggest there is a tendency for the VIX futures curve to be upward sloping. In addition, the ranges of prices for the VIX as well as the various points along the VIX futures curve and VIX term structures convey that the volatility of the prices themselves decreases as the time to maturity increases, and this is confirmed by the standard deviation readings. The intuition for this was discussed in the VIX futures pricing section. Essentially as the tenor of the futures contract rises, the uncertainty around the concavity adjustment should fall. Skewness readings are in line with implied volatilities in other asset classes. Kurtosis figures indicate fat tails for the underlying VIX, but not for the set of futures. The short-end of the resulting VIX term structure displays fat tails.

Futures						
	Near-future $\mathcal{F}_{0,1 \rightarrow 2}^{VIX}$	Next-future $\mathcal{F}_{0,2 \rightarrow 3}^{VIX}$	3rd month $\mathcal{F}_{0,3 \rightarrow 4}^{VIX}$	4th month $\mathcal{F}_{0,4 \rightarrow 5}^{VIX}$	5th month $\mathcal{F}_{0,5 \rightarrow 6}^{VIX}$	6th month $\mathcal{F}_{0,6 \rightarrow 7}^{VIX}$
Average	23.7	24.2	24.4	24.5	24.6	24.6
Max	63.7	56.9	52.8	48.5	46.4	44.4
Min	10.5	11.9	12.8	13.3	13.7	14.1
Std Dev	10.1	9.01	8.27	7.76	7.42	7.13
Skew	1.37	0.99	0.74	0.53	0.39	0.29
Kurt	2.25	1.19	0.59	0.0	-0.32	-0.49

Term structure (Constructed from futures prices)							
	VIX index $\sigma_{0 \rightarrow 1}^{VIX}$	2-month $\sigma_{0 \rightarrow 2}^{VIX}$	3-month $\sigma_{0 \rightarrow 3}^{VIX}$	4-month $\sigma_{0 \rightarrow 4}^{VIX}$	5-month $\sigma_{0 \rightarrow 5}^{VIX}$	6-month $\sigma_{0 \rightarrow 6}^{VIX}$	7-month $\sigma_{0 \rightarrow 7}^{VIX}$
Average	23.4	24.8	24.7	24.6	24.6	24.6	24.6
Max	74.3	78.7	70.6	65.7	62.1	59.3	57.1
Min	9.9	10.7	11.6	12.1	12.5	12.7	12.9
Std Dev	11.3	11.0	10.0	9.42	9.00	8.67	8.41
Skew	1.76	1.51	1.28	1.12	0.99	0.89	0.80
Kurt	4.08	3.21	2.23	1.70	1.28	0.96	0.71

Table 3

Evaluating the forecast biasness of the VIX term structure

Full sample: 4-January-2006 to 1-February-2012 (318 weekly observations)

Table contains the results for the joint hypothesis test $\alpha = 0$ and $\beta = 1$, and the individual tests $\beta = 1$ and $\beta = 0$.

$$\left(\frac{1}{k}\right) \sum_{i=1}^{k-1} \left[(\sigma_{i,i \rightarrow i+j}^{VIX})^2 - (\sigma_{0,0 \rightarrow j}^{VIX})^2 \right] = \alpha_0 + \beta_0 \left[(\sigma_{0,0 \rightarrow k+j}^{VIX})^2 - (\sigma_{0,0 \rightarrow j}^{VIX})^2 \right] + \sum_{i=1}^{k-1} u_i$$

Joint null hypothesis $\alpha = 0$ and $\beta = 1$

	7-1 [†]	6-1	5-1	4-1	3-1	2-1
α_0	-0.002	-0.002*	-0.002*	-0.002*	-0.001*	0.000*
(s.e.)	(0.010)	(0.009)	(0.008)	(0.006)	(0.004)	(0.002)
β_0	0.802	0.713*	0.593*	0.446*	0.236*	0.011*
(s.e.)	(0.117)	(0.133)	(0.147)	(0.165)	(0.187)	(0.081)
R-squared	0.301	0.242	0.174	0.107	0.037	0.0002
N	318					

Single null hypothesis $\beta = 1$

	7-1 [†]	6-1	5-1	4-1	3-1	2-1
β_0	0.802	0.713*	0.593*	0.446*	0.236*	0.011*
(s.e.)	(0.117)	(0.133)	(0.147)	(0.165)	(0.187)	(0.081)

Single null hypothesis $\beta = 0$

	7-1 [†]	6-1	5-1	4-1	3-1	2-1
β_0	0.802*	0.713*	0.593*	0.446*	0.236	0.011
(s.e.)	(0.117)	(0.133)	(0.147)	(0.165)	(0.187)	(0.081)

[†] Tests the 7-month point on the VIX term structure versus the subsequent changes in the VIX index (1-month)

* Indicates rejection at the 5% level

Table 4

Evaluating the forecast biasness of the VIX term structure

Pre-crisis period: January-2006 to December-2007 (104 weekly observations)

Table contains the results for the joint hypothesis test $\alpha = 0$ and $\beta = 1$, and the individual tests $\beta = 1$ and $\beta = 0$ for the pre-crisis sub-period.

$$\left(\frac{1}{k}\right) \sum_{i=1}^{k-1} \left[(\sigma_{i,i \rightarrow i+j}^{VIX})^2 - (\sigma_{0,0 \rightarrow j}^{VIX})^2 \right] = \alpha_0 + \beta_0 \left[(\sigma_{0,0 \rightarrow k+j}^{VIX})^2 - (\sigma_{0,0 \rightarrow j}^{VIX})^2 \right] + \sum_{i=1}^{k-1} u_i$$

Joint null hypothesis $\alpha = 0$ and $\beta = 1$

	7-1 [†]	6-1	5-1	4-1	3-1	2-1
α_0	0.002	0.002	0.001	0.003	-0.002*	0.000*
(s.e.)	(0.003)	(0.003)	(0.002)	(0.002)	(0.001)	(0.001)
β_0	0.585	0.639	0.701	0.677	0.527*	0.186*
(s.e.)	(0.226)	(0.213)	(0.179)	(0.134)	(0.110)	(0.107)
R-squared	0.089	0.116	0.152	0.171	0.139	0.052
N	104					

Single null hypothesis $\beta = 1$

	7-1 [†]	6-1	5-1	4-1	3-1	2-1
β_0	0.585	0.639	0.701	0.677*	0.527*	0.186*
(s.e.)	(0.226)	(0.213)	(0.179)	(0.134)	(0.110)	(0.107)

Single null hypothesis $\beta = 0$

	7-1 [†]	6-1	5-1	4-1	3-1	2-1
β_0	0.585*	0.639*	0.701*	0.677*	0.527*	0.186
(s.e.)	(0.226)	(0.213)	(0.179)	(0.134)	(0.110)	(0.107)

[†] Tests the 7-month point on the VIX term structure versus the subsequent changes in the VIX index (1-month)

* Indicates rejection at the 5% level

Table 5

Evaluating the forecast biasness of the VIX term structure

Crisis period: January-2008 to April-2009 (66 weekly observations)

Table contains the results for the joint hypothesis test $\alpha = 0$ and $\beta = 1$, and the individual tests $\beta = 1$ and $\beta = 0$ for the sub-period that covers the credit crisis of 2008-9.

$$\left(\frac{1}{k}\right) \sum_{i=1}^{k-1} \left[(\sigma_{i,i \rightarrow i+j}^{VIX})^2 - (\sigma_{0,0 \rightarrow j}^{VIX})^2 \right] = \alpha_0 + \beta_0 \left[(\sigma_{0,0 \rightarrow k+j}^{VIX})^2 - (\sigma_{0,0 \rightarrow j}^{VIX})^2 \right] + \sum_{i=1}^{k-1} u_i$$

Joint null hypothesis $\alpha = 0$ and $\beta = 1$

	7-1 [†]	6-1	5-1	4-1	3-1	2-1
α_0	0.040	0.032	0.024	0.016*	0.008*	0.003*
(s.e.)	(0.036)	(0.035)	(0.031)	(0.024)	(0.015)	(0.007)
β_0	1.12	0.990	0.828	0.628*	0.353*	0.075*
(s.e.)	(0.124)	(0.127)	(0.127)	(0.124)	(0.156)	(0.087)
R-squared	0.485	0.393	0.293	0.189	0.076	0.009
N	66					

Single null hypothesis $\beta = 1$

	7-1 [†]	6-1	5-1	4-1	3-1	2-1
β_0	1.12	0.990	0.828	0.628*	0.353*	0.075*
(s.e.)	(0.124)	(0.127)	(0.127)	(0.124)	(0.156)	(0.087)

Single null hypothesis $\beta = 0$

	7-1 [†]	6-1	5-1	4-1	3-1	2-1
β_0	1.12*	0.990*	0.828*	0.628*	0.353*	0.075
(s.e.)	(0.124)	(0.127)	(0.127)	(0.124)	(0.156)	(0.087)

[†] Tests the 7-month point on the VIX term structure versus the subsequent changes in the VIX index (1-month)

* Indicates rejection at the 5% level

Table 6

Evaluating the forecast biasness of the VIX term structure

Post-crisis period: April-2009 to February-2012 (149 weekly observations)

Table contains the results for the joint hypothesis test $\alpha = 0$ and $\beta = 1$, and the individual tests $\beta = 1$ and $\beta = 0$ for the sub-period following the credit crisis of 2008-9.

$$\left(\frac{1}{k}\right) \sum_{i=1}^{k-1} \left[(\sigma_{i,i \rightarrow i+j}^{VIX})^2 - (\sigma_{0,0 \rightarrow j}^{VIX})^2 \right] = \alpha_0 + \beta_0 \left[(\sigma_{0,0 \rightarrow k+j}^{VIX})^2 - (\sigma_{0,0 \rightarrow j}^{VIX})^2 \right] + \sum_{i=1}^{k-1} u_i$$

Joint null hypothesis $\alpha = 0$ and $\beta = 1$

	7-1 [†]	6-1	5-1	4-1	3-1	2-1
α_0	-0.018*	-0.015*	-0.015*	-0.008*	-0.004*	0.000*
(s.e.)	(0.008)	(0.008)	(0.007)	(0.006)	(0.004)	(0.002)
β_0	0.764*	0.620*	0.434*	0.221*	-0.003*	-0.160*
(s.e.)	(0.223)	(0.218)	(0.217)	(0.216)	(0.162)	(0.071)
R-squared	0.151	0.101	0.053	0.016	0.000	0.000
N	148					

Single null hypothesis $\beta = 1$

	7-1 [†]	6-1	5-1	4-1	3-1	2-1
β_0	0.764	0.620	0.434*	0.221*	-0.003*	-0.160*
(s.e.)	(0.223)	(0.218)	(0.217)	(0.216)	(0.162)	(0.071)

Single null hypothesis $\beta = 0$

	7-1 [†]	6-1	5-1	4-1	3-1	2-1
β_0	0.764*	0.620*	0.434*	0.221	-0.003	-0.160*
(s.e.)	(0.223)	(0.218)	(0.217)	(0.216)	(0.162)	(0.071)

[†] Tests the 7-month point on the VIX term structure versus the subsequent changes in the VIX index (1-month)

* Indicates rejection at the 5% level

Table 7

Regression results of weekly changes in realized arbitrage profits on changes in open interest, capital flow, safe haven performance, tail risk hedge costs, and credit risk factors

Table contains the regression results of weekly changes in realized arbitrage profits on changes in open interest, capital flow, safe haven performance, tail risk hedge costs, and credit risk factors. Arbitrage profits are estimated by selling the 2-month term structure and simultaneously buying the 7-month term structure at pre-defined ratios. The open interest variable describes the total open interest for VIX futures on the CBOE. Capital flows are estimated from total global hedge fund net asset values, as tracked by Hedge Fund Research Inc.. Safe have performance is proxied by the price of gold spot. Tail risk hedge costs are estimated by the relative cost of out-of-the-money puts on the Brazilian real versus the US dollar. Changes in swap spreads for 1-year tenor are used to represent changes in credit risk in the financial system.

Full sample: January-2006 to February-2012		
Explanatory variable	Regression coefficient	T-statistic
Open interest	0.039	1.16
Capital flow	1.693	2.22*
Safe haven performance	0.144	1.09
Tail risk hedge costs	0.016	0.38
Credit risk	-0.013	-0.55
F	5.94**	
R-squared	0.053	
N	317	
Post-crisis sample: April-2009 to 1-February-2012		
Explanatory variable	Regression coefficient	T-statistic
Open interest	0.112	1.95*
Capital flow	3.02	2.40*
Safe haven performance	-0.181	-1.61
Tail risk hedge costs	-0.050	-0.46
Credit risk	-0.0005	-0.02
F	18.1**	
R-squared	0.121	
N	147	

* Indicates significance at the 5% level

** Indicates significance at the 1% level

Table 8

Tests of VIX as a forecast of future realized volatility in the SPX

Table contains the results of a test I carried out using the data in this paper which assesses the forecast accuracy of the VIX. Note the main paper addresses VIX futures, this is a separate question altogether. In general, ex ante estimates overshoot ex post realized volatility in the S&P 500. Beta coefficients below one imply VIX ex ante estimates overshoot ex post realizations of S&P 500 realized volatility. The highest coefficients were recorded in 2009, a period of great financial distress. These results are consistent with Fleming (1998), Bakshi and Kapadia (2002), Buraschi and Jackwerth (2001), Coval and Shumway (2001), and Pan (2002).

Full sample: Daily data 2006-2012

$$\sigma_{t+1}^{SPX} = \beta_0 + \beta_1 \left(\sigma_t^{VIX} \right) + u_{t+1}$$

	2007	2008	2009	2010	2011	2012
β_0	0.082	0.144	0.125	0.159	0.144	0.097
β_1	0.595*	0.516*	0.783*	0.402*	0.461*	0.738*

* Indicates rejection at the 5% level