

A Tale of Two Indices

PETER CARR AND LIUREN WU

PETER CARR

is the director of the Quantitative Finance Research group at Bloomberg LP and the director of the Masters in Mathematical Finance program at the Courant Institute of New York University, NY.
pcarr4@bloomberg.com

LIUREN WU is an associate professor of economics and finance at the Zicklin School of Business, Baruch College, City University of New York, NY.
liuren_wu@baruch.cuny.edu

In 1993, the Chicago Board of Options Exchange (CBOE) introduced the CBOE Volatility Index. This index has become the de facto benchmark for stock market volatility. On September 22, 2003, the CBOE revamped the definition and calculation of the volatility index and back-calculated the new index to 1990 based on historical option prices. On March 26, 2004, the CBOE launched a new exchange, the Chicago Futures Exchange, and started trading futures on the new volatility index. Options on the new volatility index are also planned. This article describes the major differences between the old and the new volatility indexes, derives the theoretical underpinnings for the two indexes, and discusses the practical motivations behind the recent switch. It also looks at the historical behavior of the new volatility index and discusses the pricing of VIX futures and options.

In 1993, the Chicago Board of Options Exchange (CBOE) introduced the CBOE Volatility Index (VIX). This index has become the de facto benchmark for stock market volatility. The original construction of this volatility index uses options data on the S&P 100 index (OEX) to compute an average of the Black and Scholes [1973] option implied volatility with strike prices close to the current spot index level and maturities interpolated at about one month. The market often regards this implied volatility measure as a forecast of subsequent realized volatility and also as an indicator of market stress (Whaley [2000]).

On September 22, 2003, the CBOE revamped the definition and calculation of the VIX and back-calculated the new VIX to 1990

based on historical option prices. The new definition uses the S&P 500 index (SPX) to replace the OEX as the underlying stock index. Furthermore, the new index measures a weighted average of option prices across all strikes at two nearby maturities. On March 26, 2004, the CBOE launched a new exchange, the Chicago Futures Exchange (CFE), to start trading futures on the new VIX. At the time of writing, options on the VIX are also planned.

In this article, we describe the major differences in the definitions and calculations of the old and the new volatility indexes. We derive the theoretical underpinnings of the two indexes and discuss the practical motivations for the switch from the old to the new VIX. We also study the historical behavior of the new volatility index and analyze how it interacts with stock index returns and realized volatilities. Finally, we discuss how to use options on the underlying S&P 500 index to define valuation bounds on VIX futures and how to exploit information in the underlying options market and VIX futures to price options on the new VIX.

DEFINITIONS AND CALCULATIONS

The Old VXO

The CBOE renamed the old VIX the VXO, and it continues to provide quotes on this index. The VXO is based on options on the OEX. It is an average of the Black-Scholes implied volatilities on eight near-the-money

options at the two nearest maturities. When the time to the nearest maturity is within eight calendar days, the next two nearest maturities are used instead.

At each maturity, the CBOE chooses two call and two put options at the two strike prices that straddle the spot level and are nearest to it. The CBOE first averages the two implied volatilities from the put and call at each strike price and then linearly interpolates between the two average implied volatilities at the two strike prices to obtain the at-the-money spot implied volatility. The interpolated at-the-money implied volatilities at the two maturities are further interpolated along the maturity dimension to create a 22-trading-day volatility, which constitutes the VXO.

The Black-Scholes implied volatility is the annualized volatility that equates the Black-Scholes formula value to the options market quote. The annualization is based on an actual/365 day-counting convention. Instead of using this implied volatility directly, the CBOE introduced an artificial "trading-day conversion" into the calculation of the VXO. Specifically, let $ATMV(t, T)$ denote the time- t Black-Scholes at-the-money implied volatility as an annualized percentage with expiry date T . The CBOE converts this percentage to "trading-day" volatility $TV(t, T)$ as

$$TV(t, T) = ATMV(t, T) \sqrt{NC} / \sqrt{NT} \quad (1)$$

where NC and NT are the number of actual calendar days and the number of trading days between time t and the option expiry date T , respectively. The CBOE converts the number of calendar days into the number of trading days according to the following formula:

$$NT = NC - 2 \times \text{int}(NC/7) \quad (2)$$

The VXO represents an interpolated trading-day volatility at 22 trading days based on the two trading-day volatilities at the two nearest maturities ($TV(t, T_1)$ and $TV(t, T_2)$):

$$VXO_t = TV(t, T_1) \frac{NT_2 - 22}{NT_2 - NT_1} + TV(t, T_2) \frac{22 - NT_1}{NT_2 - NT_1} \quad (3)$$

where NT_1 and NT_2 denote the number of trading days between time t and the two option expiry dates T_1 and T_2 , respectively.

Since each month has around 22 trading days, the VXO represents a one-month at-the-money implied volatility estimate. Nevertheless, the trading-day conversion in Equation (1) raises the level of the VXO and makes

it no longer comparable to annualized realized volatilities computed from index returns. Thus, the VXO computation methodology has drawn criticism from both academia and industry for its artificially induced upward bias.

The New VIX

In contrast to the old VXO, which is based on near-the-money Black-Scholes implied volatilities of OEX options, the CBOE calculates the new volatility index VIX using market prices instead of implied volatilities. It also uses SPX options instead of OEX options. The general formula for the new VIX calculation at time t is

$$VIX(t, T) = \frac{2}{T-t} \sum_i \frac{\Delta K_i}{K_i^2} e^{r_t(T-t)} O_i(K_i, T) - \frac{1}{T-t} \left[\frac{F_t}{K_0} - 1 \right]^2 \quad (4)$$

where T is the common expiry date for all of the options involved in this calculation, F_t is the time- t forward index level derived from coterminous index option prices, K_i is the strike price of the i -th out-of-the-money option in the calculation, $O_i(K_i, T)$ denotes the time- t midquote price of the out-of-the-money option at strike K_i , K_0 is the first strike below the forward index level F_t , r_t denotes the time- t risk-free rate with maturity T , and ΔK_i denotes the interval between strike prices, defined as $\Delta K_i = (K_{i+1} - K_i)/2$. For notational clarity, we suppress the dependence of r_t and F_t on the maturity date T as no confusion will result.

Equation (4) uses only out-of-the-money options except at K_0 , where $O_i(K_0, T)$ represents the average of the call and put option prices at this strike. Since $K_0 \leq F_t$, the average at K_0 implies that the CBOE uses one unit of the in-the-money call at K_0 . The last term in Equation (4) represents the adjustment needed to convert this in-the-money call into an out-of-the-money put using put-call parity.

The calculation involves all available call options at strikes greater than F_t and all put options at strikes lower than F_t . The bids of these options must be strictly positive to be included. At the extreme strikes of the available options, the definition for the interval ΔK is modified as follows: ΔK for the lowest strike is the difference between the lowest strike and the next lowest strike. Likewise, ΔK for the highest strike is the difference between the highest strike and the next highest strike.

To determine the forward index level F_t , the CBOE chooses a pair of put and call options with prices that are closest to each other. Then, the forward price is derived via the put-call parity relation:

$$F_t = e^{r(T-t)}(C_t(K, T) - P_t(K, T)) + K \quad (5)$$

The CBOE uses Equation (4) to calculate $VS(t, T)$ at two of the nearest maturities of the available options, T_1 and T_2 . Then, the CBOE interpolates between $VS(t, T_1)$ and $VS(t, T_2)$ to obtain a $VS(t, T)$ estimate at 30 days to maturity. The VIX represents an annualized volatility percentage of this 30-day VS , using an actual/365 day-counting convention:

$$VIX_t = 100 \sqrt{\frac{365}{30} \left[(T_1 - t) VS(t, T_1) \frac{NC_2 - 30}{NC_2 - NC_1} + (T_2 - t) VS(t, T_2) \frac{30 - NC_1}{NC_2 - NC_1} \right]} \quad (6)$$

where NC_1 and NC_2 denote the number of actual days to expiration for the two maturities. When the nearest time to maturity is 8 days or fewer, the CBOE switches to the next-nearest maturity in order to avoid microstructure effects. The annualization in Equation (6) follows the actual/365 day-counting convention and does not suffer from the artificial upward bias incurred in the VXO calculation.

ECONOMIC AND THEORETICAL UNDERPINNINGS

The Old VXO

The VXO is essentially an average estimate of the one-month at-the-money Black-Scholes implied volatility, with an artificial upward bias induced by the trading-day conversion. Academics and practitioners often regard at-the-money implied volatility as an approximate forecast for realized volatility. However, since the Black-Scholes model assumes constant volatility, there is no direct economic motivation for regarding the at-the-money implied volatility as the realized volatility forecast beyond the Black-Scholes model context. Nevertheless, a substantial body of empirical work has found that the at-the-money Black-Scholes implied volatility is an efficient, although biased, forecast of subsequent realized volatility. Examples include Latane and Rendleman [1976], Chiras and Manaster [1978], Day and Lewis [1988], Lamoureux and Lastrapes [1993], Canina and Figlewski [1993], Fleming [1998], Christensen and Prabhala [1998], and Gwilym and Buckle [1999]. Thus, references to the VXO as a forecast of subsequent realized volatility are based more on empirical evidence than on any theoretical linkages.

Carr and Lee [2003] identify an economic interpretation for at-the-money implied volatility in a theoretical framework that goes beyond the Black-Scholes

model. They show under general market settings that the time- t -at-the-money implied volatility with expiry at time T represents an accurate approximation of the conditional risk-neutral expectation of the return volatility during the time period $[t, T]$:

$$ATMV(t, T) \cong \mathbb{E}_t^{\mathbb{Q}} [RV_{t,T}] \quad (7)$$

where $\mathbb{E}_t^{\mathbb{Q}} [\cdot]$ denotes the expectation operator under the risk-neutral measure $\mathbb{Q}$ conditional on time- t filtration $\mathcal{F}_t$, and $RV_{t,T}$ denotes the realized return volatility in annualized percentages over the time horizon $[t, T]$. Appendix A details the underlying assumptions and derivations for this approximation.

The result in Equation (7) assigns new economic meanings to the VXO, which approximates the volatility swap rate with a one-month maturity, if we readjust the upward bias induced by the trading-day conversion. Volatility swap contracts are traded actively over the counter on major currencies and some equity indexes. At maturity, the long side of the volatility swap contract receives the realized return volatility and pays a fixed volatility rate, which is the volatility swap rate. A notional dollar amount is applied to the volatility difference to convert the payoff from volatility percentage points to dollar amounts. Since the contract costs zero to enter, the fixed volatility swap rate equals the risk-neutral expected value of the realized volatility.

It is worth noting that although the at-the-money implied volatility is a good approximation of the volatility swap rate, the payoff on a volatility swap is notoriously difficult to replicate. Carr and Lee [2003] derive hedging strategies for volatility swap contracts that involve dynamic trading of both futures and options.

The New VIX

The new VIX squared approximates the conditional risk-neutral expectation of the annualized return variance over the next 30 calendar days:

$$VIX_t^2 \cong \mathbb{E}_t^{\mathbb{Q}} [RV_{t,t+30}^2] \quad (8)$$

with $RV_{t,t+30}^2 = RV_{t,t+30}^2$ denoting the annualized return variance from time t to 30 calendar days later. Hence, VIX_t^2 approximates the 30-day variance swap rate. Variance swap contracts are actively traded over the counter on major equity indexes. At maturity, the long side of the variance swap contract receives a realized variance and pays a fixed variance rate, which is the variance swap rate.

The difference between the two rates is multiplied by a notional dollar amount to convert the payoff into dollar payments. At the time of entry, the contract has zero value. Hence, by no-arbitrage, the variance swap rate equals the risk-neutral expected value of the realized variance.

Although volatility swap payoffs are difficult to replicate, variance swap payoffs can be readily replicated, up to a higher-order term. The trading strategy combines a static position in a continuum of options with a dynamic position in futures. The risk-neutral expected value of the gains from dynamic futures trading is zero. The square of the VIX is a discretized version of the initial cost of the static option position required in the replication. The theoretical relation holds under very general conditions. We can think of the VIX as the variance swap rate quoted in volatility percentage points.

To understand the replication strategy and appreciate the economic underpinnings of the new VIX, we follow Carr and Wu [2004] in decomposing the realized return variance into three components:

$$RV_{t,T} = \frac{2}{T-t} \left[\int_0^{F_t} \frac{1}{K^2} (K - S_T)^+ dK + \int_{F_t}^{\infty} \frac{1}{K^2} (S_T - K)^+ dK \right] + \frac{2}{T-t} \int_t^T \left[\frac{1}{F_{s-}} - \frac{1}{F_t} \right] dF_s - \frac{2}{T-t} \iint_t^T \left[e^x - 1 - x - \frac{x^2}{2} \right] \mu(dx, ds) \quad (9)$$

where S_t denotes the time- t spot index level, $\mathbb{R}^0$ denotes the real line excluding zero, and $\mu(dx, dt)$ is a random measure that counts the number of jumps of size $(e^x - 1)$ in the index price at time t . The decomposition in Equation (9) shows that we can replicate the return variance by the sum of 1) the payoff from a static position in a continuum of European out-of-the-money options on the underlying spot across all strike prices but at the same expiry T (first line), 2) the payoff from a dynamic trading strategy holding $[2e^{-r_f(T-s)} / (T-t)] [(1/F_{s-}) - (1/F_t)]$ futures at time s (second line), and 3) a higher-order term induced by the discontinuity in the index price dynamics (third line).

Taking expectations under the risk-neutral measure $\mathbb{Q}$ on both sides, we obtain the risk-neutral expected value of the return variance on the left-hand side. We also obtain the forward value of the sum of the start-up cost of the replicating strategy and the replication error on the right-hand side. By the martingale property, the expected value of the gains from dynamic futures trading is zero under the risk-neutral measure. With determin-

istic interest rates, we have

$$\mathbb{E}_t^{\mathbb{Q}}[RV_{t,T}] = \frac{2}{T-t} e^{r_f(T-t)} \int_0^{\infty} \frac{O_t(K, T)}{K^2} dK + \varepsilon \quad (10)$$

where ε denotes the approximation error, which is zero when the index dynamics are purely continuous and of order $O[(dF/F)^3]$ when the index can jump:

$$\varepsilon = -\frac{2}{T-t} \mathbb{E}_t^{\mathbb{Q}} \int_t^T \int_{\mathbb{R}^0} \left[e^x - 1 - x - \frac{x^2}{2} \right] \nu_s(x) dx ds \quad (11)$$

where $\nu_s(x)dxdt$ is the compensator of the jump-counting measure $\mu(dx, dt)$.

The VIX definition in Equation (4) represents a discretization of the integral in the theoretical relation in Equation (10). The extra term $(F_t/K_0 - 1)^2$ in Equation (4) is an adjustment for using a portion of the in-the-money call option at $K_0 \leq F_t$. Appendix B provides a proof for the decomposition in Equation (9) and a justification for the adjustment term in Equation (4). Therefore, the new VIX index squared has a very concrete economic interpretation. It can be regarded either as the price of a portfolio of options or as an approximation of the variance swap rate up to the discretization error and the error induced by jumps.

Practical Motivation for the Switch

The CBOE's switch from the old VXO to the new VIX was motivated by both theoretical and practical considerations. First, until very recently, the exact economic meaning of the VXO, or the at-the-money implied volatility, was not clear in any theoretical framework beyond the Black-Scholes model. It merely represents a monotonic but non-linear transformation of at-the-money option prices. In contrast, the new VIX is the price of a linear portfolio of options. The economic meaning of the new VIX is much more concrete. Second, the trading-day conversion in the VXO definition induced an artificial upward bias that has drawn criticism from both academia and industry. Third, although the VXO approximates the volatility swap rate, it remains true that volatility swaps are very difficult to replicate. In contrast, Equation (9) shows that one can readily replicate the variance swap payoffs up to a higher-order error term using a static position in a continuum of European options and a dynamic position in futures trading. Therefore, despite the popularity of the VXO as a general volatility reference index, no deriva-

tive products have been launched on the VXO index. In contrast, just a few months after the CBOE switched to the new VIX definition, it started planning to launch futures and options contracts on the new VIX. VIX futures started trading on March 26, 2004, on the CFE.

HISTORICAL BEHAVIORS

Based on historical data on daily closing option prices on the S&P 500 index and the S&P 100 index, the CBOE has back-calculated the VIX to 1990 and the VXO to 1987. For our empirical work, we choose the common sample period from January 2, 1990, to October 18, 2005, spanning 5,769 calendar days. We analyze the historical behavior of the two indexes during this sample period, with a focus on the new VIX. We also download the two stock indexes OEX and SPX and compute the realized return volatilities over the same sample period. For each day t , we compute the ex post realized volatility during the next 30 calendar days according to the following equation:

$$RVol_{t,t+30} = 100 \times \sqrt{\frac{365}{30} \sum_{j=1}^{30} \left(\ln(S_{t+j}/S_{t+j-1}) \right)^2} \quad (12)$$

where S_{t+j} denotes the index level j calendar days after day t . We follow the industry standard by computing the return squared without de-meaning the return and by annualizing the volatility according to the actual/365 day-counting convention. We analyze how the volatility indexes correlate with the index returns and return volatilities.

Summary Statistics

Exhibit 1 reports summary statistics on the levels and daily differences of the two volatility indexes (VXO and VIX), and their corresponding 30-day realized volatilities, $RVol^{SPX}$ and $RVol^{OEX}$. Since the VXO has an artificial upward bias as a result of the trading-day conversion, we also compute an adjusted index (VXOA) that scales back the conversion in the VXO: $VXOA = \sqrt{22/30} VXO$, where we approximately regard the 22 trading days as coming from 30 actual calendar days. All the volatility series are represented in percentage volatility points.

Since the VIX squared approximates the 30-day variance swap rate on the SPX and the VXOA approximates the 30-day volatility swap rate on the OEX, Jensen's inequality dictates that the VIX should be higher than the VXOA if the risk-neutral expected values of the realized volatilities on the two underlying stock indexes (OEX and SPX) are similar in magnitude:

$$VIX_t^2 \equiv \mathbb{E}_t^Q \left[\left(RVol_{t,t+30}^{SPX} \right)^2 \right] = \left(\mathbb{E}_t^Q \left[RVol_{t,t+30}^{SPX} \right] \right)^2 + \text{Var}_t^Q \left(RVol_{t,t+30}^{SPX} \right) \quad (13)$$

$$VXOA_t \equiv \mathbb{E}_t^Q \left[RVol_{t,t+30}^{OEX} \right] \quad (14)$$

$$VIX_t^2 - VXOA_t^2 \equiv \text{Var}_t^Q \left(RVol_{t,t+30}^{SPX} \right) \text{ if } \mathbb{E}_t^Q \left[RVol_{t,t+30}^{SPX} \right] \equiv \mathbb{E}_t^Q \left[RVol_{t,t+30}^{OEX} \right] \quad (15)$$

Exhibit 1 shows that the sample mean of the realized volatility on the OEX is slightly higher than that on the SPX. Nevertheless, the sample average of the VIX is higher than the sample average of the VXOA owing to Jensen's inequality. The sample average of the original VXO series is the highest, mainly because of the erroneous trading-day conversion.

EXHIBIT 1

Summary Statistics of Volatility Indexes and Realized Return Volatilities

Moments	VIX	$RVol^{SPX}$	VXO	VXOA	$RVol^{OEX}$	VIX	$RVol^{SPX}$	VXO	VXOA	$RVol^{OEX}$
	Levels					Daily Differences				
Mean	19.46	14.64	20.39	17.46	15.30	-0.00	-0.00	-0.00	-0.00	-0.00
Stdev	6.37	6.82	7.29	6.25	7.29	1.01	0.82	1.16	0.99	0.86
Skewness	0.95	1.46	0.95	0.95	1.43	0.68	0.87	0.68	0.68	0.69
Kurtosis	0.78	2.64	0.76	0.76	2.38	10.24	36.61	13.71	13.71	33.06
Auto	0.99	0.99	0.99	0.99	0.99	-0.03	0.05	-0.09	-0.09	0.06

Notes: Entries report the sample average (Mean), standard deviation (Stdev), skewness, excess kurtosis, and first-order autocorrelation (Auto) on the levels and daily differences of the new volatility index VIX, the 30-day realized volatility on SPX return ($RVol^{SPX}$), the old volatility index VXO, its bias-corrected version VXOA, and the 30-day realized volatility on OEX return ($RVol^{OEX}$). Each series has 5,769 daily observations from January 2, 1990, to October 18, 2005. All series are represented in percentage volatility points.

Comparing the volatility index with the corresponding realized volatility, we find that on average, the VIX is approximately 5 percentage points higher than the realized volatility on the SPX, and the VXOA is approximately 2 percentage points higher than the corresponding realized volatility on the OEX. To test the statistical significance of the difference between the volatility index and the realized volatility, we construct the following t -statistic:

$$t - \text{stat} = \sqrt{N} \frac{\bar{X}}{S_X} \quad (16)$$

where $N = 5,769$ denotes the number of observations, X denotes the difference between the volatility index and the realized volatility, the overline denotes the sample average, and S_X denotes the Newey and West [1987] standard deviation of X that accounts for overlapping data and serial dependence, with the number of lags optimally chosen following Andrews [1991] and an AR(1) specification. We estimate the t -statistic for $(VIX - RVol^{SPX})$ at 14.09 and for $(VIX - RVol^{OEX})$ at 6.72, both of which are highly significant.

The volatility levels show moderate positive skewness and excess kurtosis, but the excess kurtosis for daily differences is much larger, showing potential discontinuous index return volatility movements. Eraker et al. [2003] specify an index dynamics that contains constant-arrival finite-activity jumps in both the index return and the return variance rate. By estimating the model on SPX return data, they identify a strongly significant jump component in the variance rate process in addition to a significant jump component in the index return. Wu [2005] directly estimates the variance rate dynamics without specifying the return dynamics by using the VIX and various realized variance estimators constructed from tick data on SPX index futures. He also finds that the variance rate

contains a significant jump component but that the jump arrival rate is not constant over time; instead, it is proportional to the variance rate level. Furthermore, he finds that jumps in the variance rate are not rare events but arrive frequently and generate sample paths that display infinite variation.

Exhibit 2 reports the cross-correlation between the two volatility indexes (VIX_t and VXO_t) and the subsequent realized volatilities ($RVol_{t,t+30}^{SPX}$ and $RVol_{t,t+30}^{OEX}$). Each volatility index level is positively correlated with its corresponding subsequent realized volatility, but the correlation estimates become close to zero when measured in daily changes. Nevertheless, the two volatility indexes are highly correlated in both levels (0.98) and daily differences (0.86). The two realized volatility series are also highly correlated in both levels (0.99) and daily changes (0.98). Therefore, just as both stock indexes provide a general picture of the overall stock market, so both volatility indexes proxy the overall stock market volatility. Given the close correlation between the VIX and the VXO, and the planned obsolescence of the VXO, we henceforth focus our analysis on the behavior of the new VIX.

The Leverage Effect

Exhibit 3 plots the cross-correlations between SPX index returns at different leads and lags and daily changes in the volatility index VIX, with the two dash-dotted lines denoting the 95% confidence band. The instantaneous correlation estimate is strongly negative at -0.78 , but the correlation estimates at other leads and lags are much smaller. Careful inspection shows that lagged returns (within a week) show marginally significant positive correlations with daily changes in the volatility index, indicating that index returns predict future movements in the volatility index. However, index returns with

EXHIBIT 2


Cross-Correlations between Volatility Indexes and Subsequent Realized Return Volatilities

Correlation	VIX_t	$RVol_{t,t+30}^{SPX}$	VXO_t	$RVol_{t,t+30}^{OEX}$	VIX_t	$RVol_{t,t+30}^{SPX}$	VXO_t	$RVol_{t,t+30}^{OEX}$
	Levels				Daily Differences			
VIX_t	1.00	0.76	0.98	0.76	1.00	-0.04	0.86	-0.04
$RVol_{t,t+30}^{SPX}$	0.76	1.00	0.78	0.99	-0.04	1.00	-0.06	0.98
VXO_t	0.98	0.78	1.00	0.78	0.86	-0.06	1.00	-0.05
$RVOL_{t,t+30}^{OEX}$	0.76	0.99	0.78	1.00	-0.04	0.98	-0.05	1.00

Notes: Entries report the contemporaneous cross-correlation between VIX_t , SPX 30-day realized volatility ($RVol_{t,t+30}^{SPX}$), VXO_t , and OEX 30-day realized volatility ($RVol_{t,t+30}^{OEX}$) both in levels and in daily differences.

EXHIBIT 3

Cross-Correlations between Return and Volatility



Notes: The stem bars represent the cross-correlation estimates between SPX index returns at the relevant number of lags (in days) and the corresponding daily changes in VIX. The two dash-dotted lines denote the 95% confidence band.

negative lags are not significantly correlated with daily changes in the volatility index. Therefore, volatility index movements do not predict index returns.

The negative correlations between stock returns and stock return volatilities have been well documented. Nevertheless, since return volatility is not observable, the correlation can be estimated only under a structural model for return dynamics. In Exhibit 3, we use the VIX as an observable proxy for return volatility and compute the correlation across different leads and lags without resorting to a model for return dynamics. The strongly negative contemporaneous correlation between stock (index) returns and return volatilities captures the "leverage effect" first discussed by Black [1976]: given a fixed debt level, a decline in the equity level increases the leverage of the firm (market) and hence the risk for the stock (index). Various other explanations for the negative correlation have been proposed in the literature, for example, Haugen et al. [1991], Campbell and Hentschel [1992], Campbell and Kyle [1993], and Bekaert and Wu [2000].

The Federal Open Market Committee Meeting Day Effect

Balduzzi et al. [2001] find that trading volume, bid-ask spreads, and volatility on Treasury bonds and bills

increase dramatically around Federal Open Market Committee (FOMC) meeting dates. The Federal Reserve often announces changes in the Fed Funds Target Rate and its views on the overall economy during the FOMC meetings. The anticipation and ex post reaction to these announcements in monetary policy shifts and assessments create dramatic variations in trading and pricing behavior in the Treasury market. In this section, we use the VIX as a proxy for stock market volatility and investigate whether stock market volatility also shows any apparent changes around FOMC meeting days.

We download the FOMC meeting day log from Bloomberg. During our sample period, there were 144 scheduled FOMC meetings, or approximately 10 meetings per year. Exhibit 4 plots the time series of the Fed Funds Target Rates in the left panel and the basis point target changes during the scheduled FOMC meeting days in the right panel. Among the 144 meetings, 62 announced a change in the Fed Funds Target Rate. Among the 62 target moves, the change is 25 bp 45 times, 50 bp 16 times, and 75 bp once. On 25 occasions, the change is positive, representing a tightening of monetary policy, and on 37 occasions the change represents a rate cut and hence an easing of monetary policy.


Armed with the list of FOMC meeting days, we sort the VIX around the FOMC meeting days and compute the average VIX level each day from 10 days before to 10 days after each FOMC meeting day. The left panel of Exhibit 5 plots sample averages of VIX around FOMC meeting days. We observe that the average volatility level builds up before the FOMC meeting date and then drops markedly afterward. The volatility index reaches its highest level the day before the meeting and drops to the lowest level 4 days after the meeting. To investigate the significance of the drop, we measure the difference between the volatility index 1 day before and 1 day after the meeting. The mean difference is 0.6 percentage volatility points, with a *t*-statistic of 4.06.

Before the FOMC meeting, market participants disagree on whether the Fed will change the Fed Funds Target Rate, in which direction, and by how much. The fact that the option-implied stock index volatility increases prior to the meeting and drops afterward shows that the uncertainty about monetary policy has a definite impact on the volatility of the stock market. This uncertainty is resolved right after the meeting. Hence, the volatility index drops rapidly after the FOMC meeting.

Since the VIX squared can be regarded as the variance swap rate on the SPX, we also study whether the timing of a variance swap investment around FOMC

EXHIBIT 4


The Fed Funds Target Rate Changes



Notes: The solid line in the left panel plots the time series of the Fed Funds Target Rate over our sample period. The spikes in the right panel represents the target rate changes in basis points.

EXHIBIT 5

VIX Fluctuation around FOMC Meeting Days



Notes: Lines represent the sample averages of the VIX levels (left panel) and the average payoffs to long variance swap contracts, $(RV_{t,t+30} - VIX_t^2)$ (right panel), at each day within 10 days before and after the FOMC meeting days.

meeting days generates different returns. The right panel of Exhibit 5 plots the average ex post payoff from going long the swap contract around FOMC meeting days and holding the contract to maturity. The payoff is defined as the difference between the ex post realized variance and the VIX squared: $(RV_{t,t+30} - VIX_t^2)$. We find that the average

payoffs are negative by going long the swap on any day. Therefore, shorting the swap contract generates positive payoffs on average. Comparing the magnitude differences on different days, we also find that shorting the swap contract 4 days prior to the FOMC meeting day generates the highest average payoff and that shorting the variance

swap 4 days after the FOMC meeting day generates the lowest average payoff. The difference in average payoff between investments on these days is statistically significant, with a t -statistic of 9.29. Therefore, the evidence suggests that it is more profitable to short the SPX variance swap contract 4 days before an FOMC meeting than 4 days after.

Variance Risk Premia

Up to a discretization error and a jump-induced error term, the VIX squared is equal to the risk-neutral expected value of the realized variance on the SPX return during the next 30 days:

$$VIX_t^2 \cong \mathbb{E}_t^Q [RV_{t,t+30}] \quad (17)$$

We can also rewrite Equation (17) under the statistical measure $\mathbb{P}$ as

$$VIX_t^2 = \frac{\mathbb{E}_t^{\mathbb{P}}[M_{t,t+30} RV_{t,t+30}]}{\mathbb{E}_t^{\mathbb{P}}[M_{t,t+30}]} = \mathbb{E}_t^{\mathbb{P}}[RV_{t,t+30}] + \text{Cov}_t^{\mathbb{P}}\left(\frac{M_{t,t+30}}{\mathbb{E}_t^{\mathbb{P}}[M_{t,t+30}]}, RV_{t,t+30}\right) \quad (18)$$

where $M_{t,T}$ denotes a pricing kernel between times t and T . For traded assets, no-arbitrage guarantees the existence of at least one such pricing kernel (Duffie [1992]).

Equation (18) decomposes the VIX squared into two terms. The first term, $\mathbb{E}_t^{\mathbb{P}}[RV_{t,t+30}]$ represents the statistical conditional mean of the realized variance, and the second term captures the conditional covariance between the normalized pricing kernel and the realized variance. The negative of this covariance defines the time- t conditional variance risk premium (VRP):

$$VRP_t \equiv -\text{Cov}_t^{\mathbb{P}}\left(\frac{M_{t,t+30}}{\mathbb{E}_t^{\mathbb{P}}[M_{t,t+30}]}, RV_{t,t+30}\right) = \mathbb{E}_t^{\mathbb{P}}[RV_{t,t+30}] - VIX_t^2 \quad (19)$$

Taking unconditional expectations on both sides, we have

$$\mathbb{E}^{\mathbb{P}}[VRP_t] = \mathbb{E}^{\mathbb{P}}[RV_{t,t+30} - VIX_t^2] \quad (20)$$

Thus, we can estimate the average variance risk premium as the sample average of the differences between the realized return variance and the VIX squared. Over our sample period, the mean variance risk premium is estimated at -158.67 bp, with a Newey and West [1987] serial-dependence-adjusted standard error of 17.2. Hence, the mean variance risk premium is strongly negative.

Risk-averse investors normally ask for a positive risk premium for return risk. They require stock prices to appreciate by a higher percentage on average if stock

returns are riskier. In contrast, the negative variance risk premium indicates that investors require the index return variance to stay *lower* on average to compensate for higher variance risk. Therefore, whereas higher average return is regarded as compensation for higher return risk, lower average variance levels are regarded as compensation for higher variance risk. Investors are averse not only to increases in the return variance level but also to increases in the variance of the return variance.

From the perspective of a variance swap investment, the negative variance risk premium also implies that investors are willing to pay a high premium or endure an average loss when they are long variance swaps in order to receive compensation when the realized variance is high.

Dividing both sides of Equation (18) by VIX_t^2 , we can rewrite the decomposition in excess returns:

$$1 = \mathbb{E}_t^{\mathbb{P}}\left[\frac{RV_{t,t+30}}{VIX_t^2}\right] + \text{Cov}_t^{\mathbb{P}}\left(\frac{M_{t,t+30}}{\mathbb{E}_t^{\mathbb{P}}[M_{t,t+30}]}, \frac{RV_{t,t+30}}{VIX_t^2}\right) \quad (21)$$

If we regard VIX_t^2 as the forward cost of the investment in the static option position required to replicate the variance swap payoff, $(RV_{t,t+30}/VIX_t^2 - 1)$ captures the excess return from going long the variance swap. The negative of the covariance term in Equation (21) represents the conditional variance risk premium in excess return terms:


$$VRPR_t \equiv -\text{Cov}_t^{\mathbb{P}}\left(\frac{M_{t,t+30}}{\mathbb{E}_t^{\mathbb{P}}[M_{t,t+30}]}, \frac{RV_{t,t+30}}{VIX_t^2}\right) = \mathbb{E}_t^{\mathbb{P}}\left[\frac{RV_{t,t+30}}{VIX_t^2}\right] - 1 \quad (22)$$

We can estimate the mean variance risk premium in excess return form through the sample average of the realized excess returns, $ER_{t,t+30} = (RV_{t,t+30}/VIX_t^2 - 1)$, which is estimated at -40.16%, with a Newey and West [1987] standard error of 2.87%. Again, the mean variance risk premium estimate is strongly negative and highly significant. Investors are willing to endure a highly negative excess return for being long variance swaps in order to hedge away upward movements in the return variance of the stock index.

The average negative variance risk premium also suggests that shorting the 30-day variance swap and holding it to maturity generates an average excess return of 40.16%. We compute the annualized information ratio using 30-day-apart non-overlapping data, $IR = -\sqrt{12} ER/S_{ER}$, where ER denotes the time series average of the excess return and S_{ER} denotes the serial-dependence-adjusted standard deviation estimate of the excess return. The information ratio estimates average 3.52, indicating that shorting 30-day variance swaps is very profitable on average.

EXHIBIT 6

Excess Returns from Shorting 30-Day Variance Swaps



Notes: The left panel plots the time series of excess returns from shorting 30-day variance swaps on SPX and holding the contract to maturity. The right panel plots the histogram of excess returns.

To further check the historical behavior of excess returns from this investment, we plot the time series of the excess returns in the left panel and the histogram in the right panel of Exhibit 6. The time series plot shows that shorting variance swaps provides a positive return 89% of the time (5,137 out of the 5,769 daily investments). However, although the historical maximum positive return is 89.53%, the occasionally negative realizations can be as large as 242.42%. The histogram in the right panel shows that the excess return distribution is heavily negatively skewed. The high average return and high information ratio suggest that investors ask for a very high average premium to compensate for the heavily negatively skewed risk profile. The payoff from shorting variance swaps is similar to that from selling insurance, which generates a regular stream of positive premiums with small variation but with occasional exposures to large losses.

To investigate whether the classic Capital Asset Pricing Model (CAPM) can explain the risk premium from investing in variance swaps, we regress the excess returns from being long the variance swap on the excess returns from being long the market portfolio:

$$ER_{t,t+30} = \alpha + \beta (R_{t,t+30}^m - R_f) + e_t \quad (23)$$

where $(R_{t,t+30}^m - R_f)$ denotes the continuously compounded excess return to the market portfolio. If the CAPM holds,

we will obtain a highly negative beta estimate for the long variance swap return. If the CAPM can fully account for the risk premium, the estimate for the intercept α , which represents the average excess return to a market-neutral investment, will not be significantly different from zero.

We proxy the excess return to the market portfolio using the value-weighted return on all NYSE, AMEX, and NASDAQ stocks (from CRSP) minus the one-month Treasury bill rate (from Ibbotson Associates). Monthly data on the excess returns are publicly available at Kenneth French's online data library from July 1926 to September 2005. We match the sample period with our data and run the regression on monthly returns over non-overlapping data using the generalized method of moments, with the weighting matrix computed according to Newey and West [1987].

The regression estimates are as follows, with t -statistics reported in parentheses:

$$ER_t = -0.3636 - 3.7999 (R_t^m - R_f) + e_t, \quad R^2 = 19.15\% \quad (24)$$

(-65.03)(-30.10)

The beta estimate is highly negative, consistent with the general observation that index returns and volatility are negatively correlated. However, this negative beta cannot fully explain the negative premium for volatility risk. The estimate for the intercept, or the mean beta-neutral excess

return, remains strongly negative. The magnitude of α is not much smaller than the sample average of the raw excess return at -38.36% . Thus, the CAPM gets only the sign right; it cannot fully account for the large negative premium on index return variance risk. This result suggests that variability in variance constitutes a separate source of risk that the market prices heavily.

To test whether the variance risk premium is time varying, we run the following expectations-hypothesis regressions, with the t -statistics reported in parentheses:

$$\begin{aligned} RV_{t,t+30} &= -11.9006 + 0.6501VIX_t^2 + e_{t,t+30} \\ &\quad (-0.52) \quad (-4.79) \\ (RV_{t,t+30}/VIX_t^2 - 1) &= -0.4495 + 0.0001VIX_t^2 + e_{t,t+30} \\ &\quad (-14.28) \quad (1.61) \end{aligned} \quad (25)$$

Under the null hypothesis of constant variance risk premium, the first regression should generate a slope of one, and the second regression should generate a slope of zero. Zero-variance risk premium would further imply zero intercepts for both regressions. The t -statistics are computed against these null hypotheses. Since the daily series of the 30-day realized variance constitutes an overlapping series, we estimate both regressions using the generalized method of moments and construct the weighting matrix, accounting for the serial dependence according to Newey and West [1987] with 30 lags.

When the regression is run on the variance level, the slope estimate is significantly lower than the null value of

one, providing evidence that the variance risk premium VRP_t is time varying and correlated with the VIX level. When the regression is run on excess returns in the second equation, the slope estimate is no longer significantly different from zero, suggesting that the variance risk premium defined in excess return terms ($VRPR_t$) is not highly correlated with the VIX level.

Predictability of Realized Variance and Returns to Variance Swap Investments

We estimate GARCH(1,1) processes on the S&P 500 index return innovation using an AR(1) assumption on the return process. Then we compare the relative information content of the GARCH volatility and the VIX index in predicting subsequent realized return variances:

$$RV_{t,t+30} = a + b VIX_t^2 + c GARCH_t + e_{t,t+30} \quad (26)$$

where $GARCH_t$ denotes the time- t estimate of the GARCH return variance in annualized basis points. Exhibit 7 reports the generalized method of moment estimation results on restricted and unrestricted versions of this regression.

When we use either VIX^2 or $GARCH$ as the only predictor in the regression, the volatility index VIX generates an R -squared approximately 10 percentage points higher than the GARCH variance does. When we use both VIX^2 and $GARCH$ as predictors, the slope estimate on the GARCH variance is no longer statistically significant, and the R -squared is only marginally higher than using VIX^2 alone as the regressor. Thus, the GARCH variance does not provide much extra information over the VIX index.

The results in Exhibit 7 show that we can predict the realized variance using the volatility index VIX. By using variance swaps, investors can exploit such predictability and directly convert them into dollar returns. We investigate whether the predictability of return variance has been fully priced into the variance swap rate by analyzing the predictability of the excess returns from investing in a 30-day SPX variance swap and holding it to maturity.

First, we measure the monthly autocorrelation of the excess returns $ER_{t,t+30}$ using non-overlapping 30-day-apart data. The estimates average 0.12. When we run an AR(1)

EXHIBIT 7

Information Content in VIX and GARCH volatilities in predicting future realized return variances

Intercept		VIX^2		$GARCH$		R-square, %
-11.9006	(-0.52)	0.6501	(8.90)	—	—	46.87
64.0843	(3.70)	—	—	0.7456	(8.74)	35.76
-11.4691	(-0.50)	0.5873	(5.20)	0.0981	(0.88)	47.05


Notes: Entries report the estimation results on restricted and unrestricted versions of the following relation:

$$RV_{t,t+30} = a + bVIX_t^2 + cGARCH_t + e_{t,t+30}$$

The relation is estimated using the generalized method of moments. The covariance matrix is computed according to Newey and West [1987] with 30 lags. The data are daily from January 2, 1990, to October 18, 2005, generating 5,769 observations for each series.

EXHIBIT 8

Cross-Correlation between SPX Monthly Returns and Excess Returns on 30-day Variance Swaps



Notes: The stem bars represent the cross-correlation estimates between SPX returns at different lags and excess returns on investing in a 30-day variance swap and holding it to maturity. The estimates are based on monthly non-overlapping data. The two dashed lines denote the 95% confidence band. Positive numbers on the x-axis represent lags in months for index returns.

regression on the non-overlapping excess returns, the R -squared estimates average 1.58%. Thus, the predictability of excess returns through mean reversion is very low. Although the volatility level is strongly predictable, investors have priced this predictability into variance swap contracts, so that the excess returns on these swaps are not strongly predictable.

Exhibit 3 shows that SPX returns predict future movements in the VIX. Now we investigate whether we can predict the excess return on a variance swap investment using index returns. Exhibit 8 plots the cross-correlation between the excess return to the variance swap and the monthly return on SPX based on monthly sampled, and hence non-overlapping, data. The stock index return and the return on the variance swap investments show strongly negative contemporaneous correlation, but the non-overlapping series do not exhibit any significant lead-lag effects. Hence, despite the predictability in return volatilities, excess returns on variance swap investments are not strongly predictable. This result shows that the SPX options market is relatively efficient.

VIX DERIVATIVES

Given the explicit economic meaning of the new VIX and its direct link to a portfolio of options, the launch of derivatives on this index becomes the natural next step. On March 26, 2004, the CBOE launched a new exchange, the Chicago Futures Exchange, and started trading futures on the VIX. At the time of writing, options on the VIX are also being planned. In this section, we derive some interesting results regarding the pricing of VIX futures and options.

VIX Futures and Valuation Bounds

Under the assumption of no-arbitrage and continuous marking to market, the VIX futures price, F_t^{vix} , is a martingale under the risk-neutral probability measure $\mathbb{Q}$:

$$F_t^{vix} = \mathbb{E}_t^{\mathbb{Q}} [F_{T_1}^{vix}] = \mathbb{E}_t^{\mathbb{Q}} [VIX_{T_1}] \quad (27)$$

We derive valuation bounds on VIX futures that are observable from the underlying SPX options market, under two simplifying assumptions: 1) the VIX is calculated using a single strip of options maturing at $T_2 > T_1$, with $T_2 - T_1 = 30/365$, instead of two strips, and on a continuum of options prices rather than a discrete number of options; 2) the SPX index has continuous dynamics and interest rates are deterministic.

The first assumption implies that the VIX is given by

$$VIX_{T_1} = \sqrt{\frac{2}{(T_2 - T_1)B_{T_1}(T_2)} \int_0^\infty \frac{O_{T_1}(K, T_2)}{K^2} dK} \quad (28)$$

where $B_{T_1}(T_2)$ denotes the time- T_1 price of a zero bond maturing at T_2 . The second assumption further implies that the equality between the VIX squared and the risk-neutral expected value of the return variance is exact. Alternatively, we can write

$$VIX_{T_1} = \sqrt{\mathbb{E}_{T_1}^{\mathbb{Q}} RV_{T_1, T_2}} \quad (29)$$

Substituting Equation (29) in Equation (27), we have the VIX futures as

$$F_t^{vix} = \mathbb{E}_t^{\mathbb{Q}} \sqrt{\mathbb{E}_{T_1}^{\mathbb{Q}} RV_{T_1, T_2}}, \quad t \leq T_1 < T_2 \quad (30)$$

Then, the concavity of the square root and Jensen's inequality generates the following lower and upper bounds for the VIX futures:

$$\mathbb{E}_t^Q \sqrt{RV_{T_1, T_2}} \leq F_t^{vix} \leq \sqrt{\mathbb{E}_t^Q RV_{T_1, T_2}} \quad (31)$$

The lower bound is the forward volatility swap rate $L_t \equiv \mathbb{E}_t^Q \sqrt{RV_{T_1, T_2}}$, which can be approximated by a forward-starting at-the-money option. The proof is similar to that in Appendix A for the approximation of a spot volatility swap rate using the spot at-the-money option. The upper bound is the forward-starting variance swap rate quoted in volatility percentage points, $U_t \equiv \sqrt{\mathbb{E}_t^Q RV_{T_1, T_2}}$, which can be determined from the prices on a continuum of options at two maturities T_1 and T_2 :

$$\begin{aligned} U_t^2 &= \mathbb{E}_t^Q RV_{T_1, T_2} = \frac{1}{T_2 - T_1} \left[\mathbb{E}_t^Q (T_2 - t) RV_{t, T_2} - \mathbb{E}_t^Q (T_1 - t) RV_{t, T_1} \right] \\ &= \frac{2}{T_2 - T_1} \int_0^\infty \left[\frac{O_t(K, T_2)}{B_t(T_2)} - \frac{O_t(K, T_1)}{B_t(T_1)} \right] dK \end{aligned} \quad (32)$$

The width of the bounds is determined by the risk-neutral variance of the forward-starting realized volatility:

$$U_t^2 - L_t^2 = \mathbb{E}_t^Q (RV_{T_1, T_2}) - \left(\mathbb{E}_t^Q \sqrt{RV_{T_1, T_2}} \right)^2 = \text{Var}_t^Q \left(\sqrt{RV_{T_1, T_2}} \right) \quad (33)$$

When the market quote on VIX futures (F_t^{vix}) is available, we can combine it with forward-starting variance swap rates (U_t) to determine the risk-neutral variance of the future VIX:

$$\begin{aligned} \text{Var}_t^Q (VIX_{T_1}) &= \text{Var}_t^Q \left(\sqrt{\mathbb{E}_{T_1}^Q [RV_{T_1, T_2}]} \right) \\ &= \mathbb{E}_t^Q [RV_{T_1, T_2}] - \left(\mathbb{E}_t^Q \sqrt{\mathbb{E}_{T_1}^Q RV_{T_1, T_2}} \right)^2 = U_t^2 - (F_t^{vix})^2 \end{aligned} \quad (34)$$

Therefore, VIX futures provide economically relevant information not only about the future VIX level but also about the risk-neutral variance of the future VIX. We can use this information for pricing VIX options.

VIX Options

The VIX futures market, together with the SPX options market, provides the information basis for launching VIX options. To see this, we consider a call option on VIX, with the terminal payoff

$$(VIX_{T_1} - K)^+ \quad (35)$$

where K is the strike price and T_1 denotes the expiry date of the option. We have shown that we can learn the conditional risk-neutral mean (m_{1t}) and variance (m_{2t}) of VIX_{T_1} from information in the VIX futures market and the underlying SPX options market:

$$\begin{aligned} m_{1t} &= \mathbb{E}_t^Q (VIX_{T_1}) = F_t^{vix} \\ m_{2t} &= \text{Var}_t^Q (VIX_{T_1}) = U_t^2 - (F_t^{vix})^2 = \frac{2}{T_2 - T_1} \int_0^\infty \left[\frac{O_t(K, T_2)}{B_t(T_2)} - \frac{O_t(K, T_1)}{B_t(T_1)} \right] dK - (F_t^{vix})^2 \end{aligned} \quad (36)$$

Thus, under certain distributional assumptions, we can derive the value of the VIX option as a function of these two moments.

As an example, if we assume that VIX_{T_1} follows a log-normal distribution under measure $\mathbb{Q}$, we can use the Black formula to price VIX options with the two moments in Equation (36) as inputs:

$$C_t = B_t(T_1) \left[F_t^{vix} N(d_1) - KN(d_2) \right]$$

where

$$d_1 = \frac{\ln F_t^{vix}/K + \frac{1}{2} s_t^2 (T_1 - t)}{s_t \sqrt{T_1 - t}}, \quad d_2 = d_1 - s_t \sqrt{T_1 - t}$$

and s_t is the conditional annualized volatility of $\ln VIX_{T_1}$, which can be represented as a function of the first two conditional moments of VIX_{T_1} ,

$$s = \sqrt{\frac{1}{T_1 - t} \ln \frac{m_{2t} + (F_t^{vix})^2}{(F_t^{vix})^2}} \quad (37)$$

As another example, if we assume that the risk-neutral distribution of VIX_{T_1} is normal rather than lognormal, we can derive the Bachelier option pricing formula as a function of the first two observable moments of VIX_{T_1} :

$$C_t = B_t(T_1) \left[\sqrt{m_{2t}} N'(d) + (F_t^{vix} - K) N(d) \right] \quad (38)$$

with $d = (F_t^{vix} - K)/\sqrt{m_{2t}}$. For at-the-money options ($K = F_t^{vix}$), the Bachelier option pricing formula reduces to a very simple form,

$$A_t = B_t(T_1) \sqrt{m_{2t}} / \sqrt{2\pi} \quad (39)$$

CONCLUSION

The new VIX differs from the old VXO in two key aspects. First, the two indexes use different underlyings: the SPX for the new VIX versus the OEX for the old VXO. Second, the two indexes use different formulae in extracting volatility information from the options market. The new VIX is constructed from the price of a portfolio

of options and represents a model-free approximation of the 30-day return variance swap rate. The old VXO builds on the 1-month Black-Scholes at-the-money implied volatility and approximates the volatility swap rate under certain assumptions. The CBOE decided to switch from the VXO to the VIX mainly because the new VIX has a better known and more robust economic interpretation. In particular, the variance swap underlying the new VIX has a robust replicating portfolio whose option component is static. In contrast, robust replication of the volatility swap underlying the VXO index requires dynamic option trading. Furthermore, the VXO includes an upward bias induced by an erroneous trading-day conversion in its definition.

Analyzing approximately 15 years of daily data on the two volatility indices, we obtain several interesting findings on the index behavior. We find that the new VIX averages about 2 percentage points higher than the bias-corrected version of the old index, although the sample average of the 30-day realized volatility on SPX is 0.66 percentage points lower than that of OEX. The difference between the new and old volatility indexes is mainly induced by Jensen's inequality and the risk-neutral variance of realized volatility. The historical behaviors of the two volatility indexes are otherwise very similar and they move closely with each other. We also find that daily changes in the volatility indexes show very large excess kurtosis, suggesting that the volatility indexes contain large discontinuous movements.

We identify a strongly negative contemporaneous correlation between VIX and SPX index returns, confirming the "leverage effect" first documented by Black [1976]. Furthermore, although lagged index returns show marginal predictive power on the future movements of the VIX, lagged movements in the volatility index do not predict future index returns.

When we analyze VIX behavior around FOMC meeting days, during which monetary policy decisions such as Fed Funds Target Rate changes are often announced, we find that the volatility index increases prior to the FOMC meeting but drops rapidly after the meeting, showing that uncertainty about monetary policy has a direct impact on volatility in the stock market.

Since the VIX squared represents the variance swap rate on the SPX, the sample average difference between the 30-day realized return variance on the SPX and the VIX squared measures the average variance risk premium, which we estimate to be -158.67 bp and highly significant. When we represent the variance risk premium in

excess returns form, we obtain a mean estimate of -40.16% for being long a 30-day variance swap and holding it to maturity. The highly negative variance risk premium indicates that investors are averse to variations in return variance and the compensation for bearing variance risk can come in the form of a lower mean variance level under the empirical distribution than under the risk-neutral distribution.

From the perspective of variance swap investors, the negative variance risk premium indicates that investors are willing to pay a high average premium to obtain compensation (insurance) when the variance level increases. Therefore, shorting variance swaps and hence receiving the fixed leg generates positive excess returns on average. The annualized information ratio for shorting a variance swap is approximately 3.52, which is much higher than for traditional investments. Nevertheless, the excess return distribution accessed by being short variance swaps is heavily negatively skewed. Negative return realizations are few but large. The high information ratio indicates that investors ask for a high average return in order to compensate for the heavily negatively skewed risk profile. When we regress the excess returns from being long the variance swap on the stock market portfolio, we obtain a highly negative beta. However, the intercept of the regression remains highly negative, indicating that the classic Capital Asset Pricing Model cannot fully account for the negative variance risk premium. Investors regard variability in variance as a separate source of risk and charge a separate price for bearing this risk. Expectations hypothesis regressions further show that the variance risk premium in variance levels is time varying and correlated with the VIX level, but the variance risk premium in excess returns form is much less correlated with the VIX level.

We find that the VIX can predict movements in future realized variance and that GARCH volatilities do not provide extra information once the VIX is included as a regressor. Nevertheless, the strong predictability of the realized variance does not transfer to strong predictability in excess returns for investing in variance swaps.

Finally, we show that the SPX options market provides information on valuation bounds for VIX futures. The width of the bounds are determined by the risk-neutral variance for forward-starting return volatility. Furthermore, VIX futures quotes not only provide information about the risk-neutral mean of future VIX levels but also combine with information from the SPX

options market to reveal the risk-neutral variance of the VIX. This information can be used to price VIX options.

APPENDIX A

Approximating Volatility Swap Rates with At-the-Money Implied Volatilities

Let $(\Omega, \mathcal{F}, \mathbb{Q})$ be a probability space defined on a risk-neutral measure $\mathbb{Q}$. As in Carr and Lee [2003], we assume continuous dynamics for the index futures F_t under measure $\mathbb{Q}$:

$$dF_t/F_t = \sigma_t dW_t \quad (\text{A-1})$$

where the diffusion volatility σ_t can be stochastic, but its variation is assumed to be independent of the Brownian motion W_t in the price. Under these assumptions, Hull and White [1987] show that the value of a call option can be written as the risk-neutral expected value of the Black-Scholes formula, evaluated at the realized volatility. The time- t value of the at-the-money forward ($K = F_t$) option maturing at time T can be written as

$$\begin{aligned} ATMC_{t,T} = \mathbb{E}_t^{\mathbb{Q}} \left\{ F_t \left[N \left(\frac{RVol_{t,T} \sqrt{T-t}}{2} \right) - N \left(\frac{-RVol_{t,T} \sqrt{T-t}}{2} \right) \right] \right\} \end{aligned} \quad (\text{A-2})$$

where $RVol_{t,T}$ is the annualized realized return volatility over $[t, T]$

$$RVol_{t,T} \equiv \sqrt{\frac{1}{T-t} \int_t^T \sigma_s^2 ds} \quad (\text{A-3})$$

Brenner and Subrahmanyam [1988] show that a Taylor series expansion of each normal distribution function about zero implies

$$\begin{aligned} N \left(\frac{RVol_{t,T} \sqrt{T-t}}{2} \right) - N \left(\frac{-RVol_{t,T} \sqrt{T-t}}{2} \right) \\ = \frac{RVol_{t,T} \sqrt{T-t}}{\sqrt{2\pi}} + O((T-t)^{\frac{3}{2}}) \end{aligned} \quad (\text{A-4})$$

Substituting Equation (A-4) into Equation (A-2) implies that

$$ATMC_{t,T} \approx \mathbb{E}_t^{\mathbb{Q}} \left[\frac{F_t}{\sqrt{2\pi}} RVol_{t,T} \sqrt{T-t} \right] \quad (\text{A-5})$$

and hence the volatility swap rate is given by

$$\mathbb{E}_t^{\mathbb{Q}} [RVol_{t,T}] = \frac{\sqrt{2\pi}}{F_t \sqrt{T-t}} ATMC_{t,T} + O((T-t)^{\frac{3}{2}}) \quad (\text{A-6})$$

Since an at-the-money call value is concave in volatility, $(\sqrt{2\pi}/F_t \sqrt{T-t}) ATMC_{t,T}$ is a slightly downward-biased approximation of the volatility swap rate. As a result, the error term is positive. However, Brenner and Subrahmanyam show that the at-the-money implied volatility is also given by

$$ATMC_{t,T} = \frac{\sqrt{2\pi}}{F_t \sqrt{T-t}} ATMV_{t,T} + O((T-t)^{\frac{3}{2}}) \quad (\text{A-7})$$

Once again, $(\sqrt{2\pi}/F_t \sqrt{T-t}) ATMV_{t,T}$ is a slightly downward-biased approximation of the at-the-money implied volatility. Subtracting Equation (A-7) from Equation (A-6) implies that the volatility swap rate is approximated by the at-the-money implied volatility:

$$\mathbb{E}_t^{\mathbb{Q}} [RVol_{t,T}] = ATMV_{t,T} + O((T-t)^{\frac{3}{2}}) \quad (\text{A-8})$$

The leading source of error in Equation (A-6) is partially canceled by the leading source of error in Equation (A-7). As a result, this approximation has been found to be very accurate.

APPENDIX B

Replicating Variance Swaps with Options

The interpretation of the new VIX as an approximation of the 30-day variance swap rate can be derived under a much more general setting for the $\mathbb{Q}$ -dynamics of SPX index futures:

$$dF_t/F_t = \sigma_{t-} dW_t + \int_{\mathbb{R}^0} (e^x - 1) [\mu(dx, dt) - \nu_t(x) dx dt] \quad (\text{A-9})$$

where F_{t-} denotes the futures price at time t just prior to a jump, $\mathbb{R}^0$ denotes the real line excluding zero, and the random measure $\mu(dx, dt)$ counts the number of jumps of size $(e^x - 1)$ in the index futures at time t . The process $\{\nu_t(x), x \in \mathbb{R}^0\}$ compensates the jump process $J_t \equiv \int_0^t \int_{\mathbb{R}^0} (e^x - 1) \mu(dx, ds)$, so that the last term in Equation (A-9) is the increment of a $\mathbb{Q}$ -pure jump martingale. To avoid notational complexity, we assume that the jump component in the price process exhibits finite variation:

$$\int_{\mathbb{R}^0} (|x| \wedge 1) \nu_t(x) dx < \infty$$

By adding the time subscripts to σ_{t-} and $\nu_t(x)$, we allow both to be stochastic and predictable with respect to the filtration $\mathcal{F}_t$. To satisfy limited liability, we further assume the two stochastic

processes to be such that the futures price F_t is always non-negative and absorbing at the origin. Finally, with little loss of generality, we assume deterministic interest rates and dividend yields. Under these assumptions, the annualized quadratic variation on the futures return over horizon $[t, T]$ can be written as

$$RV_{t,T} = \frac{1}{T-t} \left[\int_t^T \sigma_{t-}^2 dt + \int_0^T \int_{\mathbb{R}^0} x^2 \mu(dx, dt) \right] \quad (\text{A-10})$$

Applying Itô's lemma to the function $f(F) = \ln F$, we have

$$\ln(F_T) = \ln(F_t) + \int_t^T \frac{1}{F_s} dF_s - \frac{1}{2} \int_t^T \sigma_{s-}^2 ds + \int_t^T \int_{\mathbb{R}^0} \left[x - e^x + 1 \right] \mu(dx, ds)$$

Adding and subtracting $2[(F_T/F_t) - 1] + \int_t^T x^2 \mu(dx, dt)$ and rearranging, we obtain a representation for the quadratic variation:

$$(T-t)RV_{t,T} = 2 \left[\frac{F_T}{F_t} - 1 - \ln \left(\frac{F_T}{F_t} \right) \right] + 2 \int_t^T \left[\frac{1}{F_{s-}} - \frac{1}{F_t} \right] dF_s - 2 \int_t^T \int_{\mathbb{R}^0} \left[e^x - 1 - x - \frac{x^2}{2} \right] \mu(dx, ds) \quad (\text{A-11})$$

A Taylor expansion with the remainder of $\ln F_T$ about the point F_t implies

$$\ln F_T = \ln F_t + \frac{1}{F_t} (F_T - F_t) - \int_0^1 \frac{1}{K^2} (K - F_T)^+ dK - \int_{F_t}^\infty \frac{1}{K^2} (F_T - K)^+ dK \quad (\text{A-12})$$

Plugging Equation (A-12) into Equation (A-11), we have

$$(T-t)RV_{t,T} = 2 \left[\int_0^{F_t} \frac{1}{K^2} (K - F_T)^+ dK + \int_{F_t}^\infty \frac{1}{K^2} (F_T - K)^+ dK \right] + 2 \int_t^T \left[\frac{1}{F_{s-}} - \frac{1}{F_t} \right] dF_s + 2 \int_t^T \int_{\mathbb{R}^0} \left[e^x - 1 - x - \frac{x^2}{2} \right] \mu(dx, ds) \quad (\text{A-13})$$

which is the decomposition in Equation (9) that also represents a replicating strategy for the return quadratic variation.

Taking expectations under measure $\mathbb{Q}$, we obtain the risk-neutral expected value of the return variance on the left-hand side and the cost of the replication strategy on the right-hand side:

$$\mathbb{E}_t^{\mathbb{Q}}[RV_{t,T}] = \frac{2e^{r_t(T-t)}}{T-t} \int_0^\infty \frac{O_t(K, T)}{K^2} dK - 2\mathbb{E}_t^{\mathbb{Q}} \int_t^T \int_{\mathbb{R}^0} \left[e^x - 1 - x - \frac{x^2}{2} \right] v_s(x) dx ds$$

where the first term denotes the initial cost of the static portfolio of out-of-the-money options and the second term is a higher-order error term induced by jumps.

The VIX definition in Equation (4) represents a discretization of the option portfolio. The extra term $(F_t/K_0 - 1)^2$ in the VIX definition adjusts for the in-the-money call option used at $K_0 \leq F_t$. To convert the in-the-money call option into the out-of-the-money put option, we use the put-call parity:

$$e^{r_t(T-t)} C_t(K_0, K) = e^{r_t(T-t)} P_t(K_0, T) + F_t - K_0 \quad (\text{A-14})$$

If we plug this equality into Equation (4) to convert all option prices into out-of-the-money option prices, we have

$$VS(t, T) = \frac{2}{T-t} \sum \frac{\Delta K_i}{K_i^2} e^{r_t(T-t)} O_t(K_i, T) + \frac{\Delta K_0}{(T-t)K_0^2} (F_t - K_0) - \frac{1}{T-t} \left[\frac{F_t}{K_0} - 1 \right]^2 \quad (\text{A-15})$$

where the second term on the right-hand side of Equation (A-15) is due to the substitution of the in-the-money call option at K_0 by the out-of-the-money put option at the same strike K_0 . If we further assume that the forward level is in the middle of the two adjacent strike prices and approximate the interval ΔK_0 by $F_t - K_0$, the last two terms in Equation (A-15) cancel out to obtain

$$VS(t, T) = \frac{2}{T-t} \sum \frac{\Delta K_i}{K_i^2} e^{r_t(T-t)} O_t(K_i, T) \quad (\text{A-16})$$

Thus, the VIX definition matches the theoretical relation for the risk-neutral expected value of the return quadratic variation up to a jump-induced error term and errors induced by discretization of strikes.

ENDNOTE

The authors thank Stephen Figlewski (the editor), Robert Engle, Harvey Stein, Benjamin Wurzbürger, and seminar participants at New York University and the 20th Annual Risk Management Conference in Florida for many insightful comments. All remaining errors are ours.

REFERENCES

- Andrews, D. "Heteroskedasticity and Autocorrelation Consistent Covariance Matrix Estimation." *Econometrica*, 59 (1991), pp. 817-858.
- Balduzzi, P., E.J. Elton, and T.C. Green. "Economic News and Bond Prices: Evidence from the U.S. Treasury Market." *Journal of Financial and Quantitative Analysis*, 36 (2001), pp. 523-543.
- Bekaert, G., and G. Wu. "Asymmetric Volatilities and Risk in Equity Markets." *Review of Financial Studies*, 13 (2000), pp. 1-42.

- Black, F. "Studies of Stock Price Volatility Changes." In *Proceedings of the 1976 American Statistical Association, Business and Economical Statistics Section*. Alexandria, VA: American Statistical Association, 1976, pp. 177-181.
- Black, F., and M. Scholes. "The Pricing of Options and Corporate Liabilities." *Journal of Political Economy*, 81 (1973), pp. 637-654.
- Brenner, M., and M. Subrahmanyam. "A Simple Formula to Compute the Implied Standard Deviation." *Financial Analysts Journal*, 44 (1988), pp. 80-83.
- Campbell, J.Y., and L. Hentschel. "No News Is Good News: An Asymmetric Model of Changing Volatility in Stock Returns." *Review of Economic Studies*, 31 (1992), pp. 281-318.
- Campbell, J.Y., and A.S. Kyle. "Smart Money, Noise Trading and Stock Price Behavior." *Review of Economic Studies*, 60 (1993), pp. 1-34.
- Canina, L., and S. Figlewski. "The Information Content of Implied Volatility." *Review of Financial Studies*, 6 (1993), pp. 659-681.
- Carr, P., and R. Lee. "At-the-Money Implied as a Robust Approximation of the Volatility Swap Rate." Working paper, New York University, 2003.
- Carr, P., and L. Wu. "Variance Risk Premia." Working paper, New York University and Baruch College, 2004.
- Chiras, D., and S. Manaster. "The Information Content of Option Prices and a Test of Market Efficiency." *Journal of Financial Economics*, 6 (1978), pp. 213-234.
- Christensen, B.J., and N.R. Prabhala. "The Relation between Implied and Realized Volatility." *Journal of Financial Economics*, 50 (1998), pp. 125-150.
- Day, T.E., and C.M. Lewis. "The Behavior of the Volatility Implicit in Option Prices." *Journal of Financial Economics*, 22 (1988), pp. 103-122.
- Duffe, D. *Dynamic Asset Pricing Theory*, 2nd ed. Princeton, NJ: Princeton University Press, 1992.
- Eraker, B., M. Johannes, and N. Polson. "The Impact of Jumps in Equity Index Volatility and Returns." *Journal of Finance*, 58 (2003), pp. 1269-1300.
- Fleming, J. "The Quality of Market Volatility Forecast Implied by S&P 500 Index Option Prices." *Journal of Empirical Finance*, 5 (1998), pp. 317-345.
- Gwilym, O.A., and M. Buckle. "Volatility Forecasting in the Framework of the Option Expiry Cycle." *European Journal of Finance*, 5 (1999), pp. 73-94.
- Haugen, R.A., E. Talmor, and W.N. Torous. "The Effect of Volatility Changes on the Level of Stock Prices and Subsequent Expected Returns." *Journal of Finance*, 46 (1991), pp. 985-1007.
- Hull, J., and A. White. "The Pricing of Options on Assets with Stochastic Volatilities." *Journal of Finance*, 42 (1987), pp. 281-300.
- Lamoureux, C.G., and W.D. Lastrapes. "Forecasting Stock-Return Variance: Toward an Understanding of Stochastic Implied Volatilities." *Review of Financial Studies*, 6 (1993), pp. 293-326.
- Latane, H.A., and R.J. Rendleman. "Standard Deviation of Stock Price Ratios Implied in Option Prices." *Journal of Finance*, 31 (1976), pp. 369-381.
- Newey, W.K., and K.D. West. "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix." *Econometrica*, 55 (1987), pp. 703-708.
- Whaley, R.E. "The Investor Fear Gauge." *Journal of Portfolio Management*, 26 (2000), pp. 12-17.
- Wu, L. "Variance Dynamics: Joint Evidence from Options and High-Frequency Returns." Working paper, Baruch College, 2005.

To order reprints of this article, please contact Dewey Palmieri at dpalmieri@ijjournals.com or 212-224-3675.

©Euromoney Institutional Investor PLC. This material must be used for the customer's internal business use only and a maximum of ten (10) hard copy print-outs may be made. No further copying or transmission of this material is allowed without the express permission of Euromoney Institutional Investor PLC.

The most recent two editions of this title are only ever available at <http://www.euromoneyplc.com>