A TALE OF TWO INDICES

Peter Carr
Liuren Wu

A Tale of Two Indices*

PETER CARR†

Bloomberg L.P. and Courant Institute

LIUREN WU‡

Zicklin School of Business, Baruch College

First draft: October 20, 2003

This version: March 19, 2004

*We thank Harvey Stein, and Benjamin Wurzburger for inspiring discussions. We welcome comments, including references we have inadvertently missed.

†499 Park Avenue, New York, NY 10022; tel: (212) 893-5056; fax: (917) 369-5629; pcarr4@bloomberg.com.

‡One Bernard Baruch Way, Box B10-254, New York, NY 10010; tel: (646) 312-3509; fax: (646) 312-3451; Liuren_Wu@baruch.cuny.edu; http://faculty.baruch.cuny.edu/lwu.

A Tale of Two Indices

ABSTRACT

In 1993, the Chicago Board of Options Exchange (CBOE) introduced the COBE Volatility Index (VIX). 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 VIX, and back-calculated the new VIX up to 1990 based on historical option prices. The CBOE is also planning to launch futures and options on the new VIX. In this paper, we describe the major differences between the old and the new VIXs, derive the theoretical underpinnings for the two indices, and discuss the practical motivation for the recent switch. We also study the historical behaviors of the two indices.

A Tale of Two Indices

In 1993, the Chicago Board of Options Exchange (CBOE) introduced the COBE Volatility Index (VIX). This index has become the de factor benchmark for stock market volatility. It is widely followed and has been cited in hundreds of news articles in the Wall Street Journal, Barron's and other leading financial publications.

The volatility index uses options data on S&P 100 index (OEX) and computes 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 on market stress.

On September 22, 2003, following suggestions from the industry,[1] CBOE revamped the definition and calculation of the VIX, and back-calculated the new VIX up to 1990 based on historical option prices. The new definition uses the more actively traded S&P 500 index options to replace the S&P 100 index as the underlying index. Furthermore, the new index measures a weighted average of option prices across all strikes at two nearby maturities.

Currently, the CBOE keeps track of both volatility indexes and rename the old index as VXO. The CBOE has also been planning to launch futures and options on the new VIX. In this paper, we describe the major differences in the definition and calculation of the two volatility indices. We also derive the theoretical underpinnings for the two definitions and discuss the practical motivations for the switch from the old to the new VIX. Finally, we study the historical behavior of the two volatility indexes.

1. Definitions and Calculations

1.1. The old VXO

The CBOE renames the old VIX now as VXO and continues to provide the quotes on this index. The calculation of the VXO index is based on options on the S&P 100 index (OEX). It is an average of the

[1] See "Developing the New VIX — A Practitioner's Tale," by Sandy Rattray at Goldman, Sachs & Co.

Black-Scholes implied volatility quotes on eight near-the-money options at the two nearby maturities. At each maturity, the CBOE picks the two call and two put options that are closest to the money and average their implied volatility quotes to obtain an estimate of the approximately at-the-money implied volatility at that maturity. Then, the CBOE linearly interpolates between the two at-the-money implied volatility estimates to obtain an at-the-money implied volatility estimate at the one-month maturity level. The interpolation is based on the number of business days.

1.2. The new VIX

The CBOE calculates the new volatility index, VIX, using market prices on the S&P 500 index options. The general formula for the new VIX calculation is

$$\sigma^2 = \frac{2}{T}\sum_i \frac{\Delta K}{K_i^2} e^{rT} P(K_i, T) - \frac{1}{T}\left[\frac{F}{K_0} - 1\right]^2, \tag{1}$$

where T is the common time to maturity for the all the options involved in this calculation, F is the forward index level derived from the index option prices, K_i is the strike price of the i-th out-of-the-money option in the calculation, $P(K_i,T)$ denotes the 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, r denotes the riskfree rate of maturity T, and ΔK_i denotes the interval between strike prices, defined as

$$\Delta K_i = \frac{K_{i+1} - K_i}{2}. \tag{2}$$

The formula in equation (1) only uses out-of-the-money options. Thus, $P(K_i,T)$ represents the call option price when $K_i > F$ and the put option price when $K_i < F$. When $K_i = K_0$, CBOE uses the average of the call and put option prices at this strike as the input for $P(K_0,T)$. Since $K_0 \leq F$, 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 (1) represents the adjustment term via the put-call parity to convert this in-the-money call into an out-of-money put.

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

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

The CBOE uses equation (1) to calculate σ^2 at two of the most shortest maturities of the available options, T_1 and T_2. Then, the CBOE linearly interpolates between the two σ^2 to obtain a σ^2 at 30-day maturity. The VIX represents the annualized percentage of this 30-day σ,

$$VIX = 100\sqrt{\frac{365}{30}\left[T_1\sigma_1^2\frac{N_{T_2}-30}{N_{T_2}-N_{T_1}} + T_2\sigma_2^2\frac{30-N_{T_1}}{N_{T_2}-N_{T_1}}\right]}, \tag{3}$$

where N_{T_1} and N_{T_2} denote the number of actual days to expirations for the two maturities. When the shortest maturity falls within eight days, the CBOE switches to another maturity to avoid microstructure effects at very short option maturities.

2. Theoretical Underpinnings

2.1. The old VXO

The VXO is essentially an average estimate of the one-month at-the-money Back-Scholes implied volatility. Both the academics and practitioners often regard the at-the-money implied volatility as an approximate forecast for realized volatility. However, since the Black-Scholes model assumes deterministic volatility, there is no direct economic motivation for regarding the at-the-money implied volatility as the realized volatility forecast. 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 of such studies include Latane and Rendleman (1976), Chiras and Manaster (1978), Day and Lewis (1988), Day and Lewis (1992), Lamoureux and Lastrapes (1993), Canina and Figlewski (1993), Day and Lewis (1994), Jorion (1995), Fleming (1998), Christensen and Prabhala (1998), Gwilym and Buckle (1999), Hol and Koopman (2000), Blair, Poon, and

Taylor (2000), Hansen (2001), Christensen and Hansen (2002), Tabak, Chang, and de Andrade (2002), Shu and Zhang (2003), and Neely (2003).

Thus, the wide reference to the VXO is more based on empirical evidence on its relevance to realized volatility than based on theoretical linkages unless under the very strict assumption of the Black-Scholes model. This situation changed recently when Carr and Lee (2003) show that at-the-money implied volatility represents an accurate approximation of the conditional risk-neutral expectation of the return volatility under general market settings. Their result indicates that the at-the-money implied volatility actually approximates the volatility swap rate. 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 rate. Since the contract costs zero to enter, the fixed volatility swap rate equals the risk-neutral expected value of the realized volatility.

Carr and Lee (2003) assume continuous futures price $\mathbb{Q}$-dynamics as follows,

$$dF_t/F_t = \sigma_t dW_t, \tag{4}$$

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 is just the risk-neutral expected value of the Black Scholes formula value, considered as a function of the random realized volatility. In the special case when the call is at-the-money ($K = F$), we have the time-0 value of the call option maturing at time T as

$$ATMC_0 = \mathbb{E}_0^{\mathbb{Q}} \left\{ F_0 \left[N\left(\frac{\sigma_T \sqrt{T}}{2} \right) - N\left(-\frac{\sigma_T \sqrt{T}}{2} \right) \right] \right\}, \tag{5}$$

where σ_T is the random volatility realized over $[0, T]$:

$$\sigma_T \equiv \sqrt{\frac{\langle X \rangle_T}{T}}. \tag{6}$$

As first shown in Brenner and Subrahmanyam (1988), a Taylor series expansion of each normal distribution function about zero implies:

$$N\left(\frac{\sigma\sqrt{T}}{2}\right) - N\left(-\frac{\sigma\sqrt{T}}{2}\right) = \frac{\sigma\sqrt{T}}{\sqrt{2\pi}} + O(T^{\frac{3}{2}}). \tag{7}$$

Substituting (7) in (5) implies that:

$$ATMC_0 \approx \mathbb{E}_0^{\mathbb{Q}}\left[\frac{F_0}{\sqrt{2\pi}}\sigma_T\sqrt{T}\right], \tag{8}$$

and hence the volatility swap rate is given by:

$$\mathbb{E}_0^{\mathbb{Q}}\sigma_T = \frac{\sqrt{2\pi}}{F_0\sqrt{T}}ATMC_0 + O(T^{\frac{3}{2}}). \tag{9}$$

Since an at-the-money call value is concave in volatility, $\frac{\sqrt{2\pi}}{F_0\sqrt{T}}ATMC_0$ is a slightly downward biased approximation of the volatility swap rate. As a result, the coefficient on $T^{\frac{3}{2}}$ is positive. However, Brenner and Subrahmanyam have shown that the at-the-money implied volatility (ATMV) is also given by:

$$ATMV = \frac{\sqrt{2\pi}}{F_0\sqrt{T}}ATMC_0 + O(T^{\frac{3}{2}}). \tag{10}$$

Once again, $\frac{\sqrt{2\pi}}{F_0\sqrt{T}}ATMC_0$ is a slightly downward biased approximation of the at-the-money implied volatility and hence the coefficient on $T^{\frac{3}{2}}$ is positive. Subtracting equation (10) from (10) implies that the initial volatility swap rate is approximated by the initial at-the-money implied volatility:

$$\mathbb{E}_0^{\mathbb{Q}}\sigma_T = ATMV + O(T^{\frac{3}{2}}). \tag{11}$$

In fact, the leading source of error in (9) is partially cancelled by the leading source of error in (10). As a result, this approximation has been found to be extremely accurate. The shorter the time to maturity, the better the approximation. Both concepts coincide with each other and instantaneous volatility as time to maturity goes to zero.

It is worth noting that although the at-the-money implied volatility happens to approximate well the volatility swap rate, the volatility swap contract itself is notoriously difficult to replicate and hedge.

Carr and Lee (2003) have derived some results on this, but the hedging strategies that propose are by any means complicated and difficult (costly) to implement in practice.

2.2. The new VIX

In contrast to the old VXO, the new VIX underlies the annualized conditional return quadratic variation under the risk-neutral measure, and hence the 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 figure 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, an sample estimate of the return quadratic variation. Although volatility swap contracts are difficult to hedge, the variance swap contracts can be readily replicated, up to a higher-order term, using a static position on a continuum of options and a dynamic position on futures trading. The risk-neutral expected value of the dynamic futures trading is zero. The VIX calculation represents a discretized version of the initial cost of the the continuum of options in the replication. The theoretical relation holds under very general conditions.

Formally, we use S_t to denote the spot price of an asset at time $t \in [0, \mathcal{T}]$, where $\mathcal{T}$ is some arbitrarily distant horizon. We use F_t to denote the time t price of a futures contract with maturity $T > t$ that marks to market continuously. No arbitrage implies that there exists a risk-neutral probability measure $\mathbb{Q}$ defined on a probability space $(\Omega, \mathcal{F}, \mathbb{Q})$ such that the futures price F_t solves the following stochastic differential equation,

$$dF_t/F_{t-} = \sigma_{t-} dW_t + \int_{\mathbb{R}^0} (e^x - 1) \left[\mu(dx, dt) - \nu_t(x) dx dt\right], \quad t \in [0, \mathcal{T}], \tag{12}$$

starting at some fixed and known value $F_0 > 0$. In equation (12), F_{t-} denotes the futures price at time t just prior to a jump, $\mathbb{R}^0$ denotes the real line excluding zero, W_t is a $\mathbb{Q}$ standard Brownian motion, and the random measure $\mu(dx, dt)$ counts the number of jumps of size e^x in the asset price at time t. The process $\{\nu_t(x), x \in \mathbb{R}^0, t \in [0, \mathcal{T}]\}$ 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 (12) is the increment of a $\mathbb{Q}$-pure jump martingale. The process $\nu_t(x)$ must have the following properties (see Prokhorov and Shiryaev (1998)),

$$\nu_0(x) = 0, \quad \nu_t(0) = 0, \quad \int_{\mathbb{R}^0} \left(|x|^2 \wedge 1\right) \nu_t(x) dx < \infty, \quad t \in [0, \mathcal{T}].$$

The literature often refers to $\nu_t(x)$ as the *compensator* or the *local density* of the jumps. Thus, equation (12) models the futures price change as the sum two orthogonal martingale components: a purely continuous martingale and a purely discontinuous (jump) martingale. This decomposition is generic for a martingale (Jacod and Shiryaev (1987), page 84).

To avoid notational complexity, we assume that the jump component in the price process exhibits finite variation,

$$\int_{\mathbb{R}^0} (|x| \wedge 1) \nu_t(x) dx < \infty, \quad t \in [0, \mathcal{T}].$$

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 nonnegative and absorbing at the origin. Finally, with little loss of generality, we assume constant interest rates and dividend yields. Under this assumption, the futures price and the forward price are identical.

Under the specification in (12), the quadratic variation on the futures return over horizon T can be written as

$$[\ln F_T, \ln F_T] = \int_0^T \sigma_{t_-}^2 dt + \int_0^T \int_{\mathbb{R}^0} x^2 \mu(dx, dt). \tag{13}$$

Under this general setting, Carr and Wu (2003) show that the time-0 risk-neutral expected value of the quadratic return variation over horizon T defined in (13) can be approximated by the value of a continuum of European out-of-the-money options across all strikes and maturing all at time T,

$$\mathbb{E}_0^{\mathbb{Q}} [\ln F_T, \ln F_T] = e^{rT} \int_0^\infty \frac{2P_0(K,T)}{K^2} dK + \varepsilon, \tag{14}$$

where ε denotes the approximation error and $P_0(K,T)$ denotes the time-0 value of an out-of-the-money option with strike price K and expiring at time T (a call option when $K > F_0$ and a put option when $K \leq F_0$). The approximation error ε is zero when the futures dynamics is purely continuous. When the futures dynamics has a discontinuous component, the approximation error ε is of order $O\left[\left(\frac{dF}{F}\right)^3\right]$ and is determined by the compensator of this discontinuous component,

$$\varepsilon = -2\mathbb{E}_0^{\mathbb{Q}} \int_0^T \int_{\mathbb{R}^0} \left[e^x - 1 - x - \frac{x^2}{2} \right] \nu_t(x) dx dt. \tag{15}$$

We refer the interested readers to Appendix A for the details of the proof. Carr and Madan (1998) and Demeterfi, Derman, Kamal, and Zou (1999a,b) have derived similar relations under the assumption of continuous sample path for the underlying futures.

It is important to note that the return quadratic variation can be written as

$$\begin{aligned} [\ln F_T, \ln F_T] &= 2\left[\int_0^{F_0} \frac{1}{K^2}(K - S_T)^+ dK + \int_{F_0}^{\infty} \frac{1}{K^2}(S_T - K)^+ dK\right] \\ &\quad + 2\int_0^T \left[\frac{1}{F_{s-}} - \frac{1}{F_0}\right] dF_s \\ &\quad - 2\int_0^T \int_{\mathbb{R}^0} \left[e^x - 1 - x - \frac{x^2}{2}\right] \mu(dx, ds). \end{aligned} \tag{16}$$

Thus, we can replicate the return quadratic variation up to time T by the sum of (i) the payoff from a static position in $\frac{dK}{K^2}$ European options on the underlying spot at strike K and expiry T (first line), (ii) the payoff from a dynamic trading strategy holding $2e^{-r(T-s)}\left[\frac{1}{F_{s-}} - \frac{1}{F_t}\right]$ futures at time s (second line), and (iii) a higher-order error term induced by the discontinuity in the futures price dynamics (third line). The options are all out-of-the money forward, i.e., call options when $F_t > K$ and put options when $K \leq F_t$.

Taking expectations under measure $\mathbb{Q}$ on both sides, we obtain the risk-neutral expected value of the quadratic variation on the left hand side. We also obtain the forward value of the sum of the startup 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.

Carr and Wu (2003) show that under commonly used jump-diffusion stochastic volatility models and reasonable parameters, the replication error term ε is small and negligible. The CBOE's calculation of the new VIX in equation (1) represents a discretization of the integral in equation (14) and therefore a model-free approximation of the annualized conditional quadratic variation of the index return.

Comparing VIX's definition in equation (1) to the theoretical relation in (13), we observe an extra term in the VIX's definition, $(F/K_0 - 1)^2$. This term is zero when $F = K_0$. Under normally conditions, $F \geq K_0$ because the CBOE set K_0 equal to the first strike price available that is below the forward value. Furthermore, instead of using all out-of-money options, the CBOE uses the average of the call and put option price at strike K_0. At stike $K_0 \leq F$, the put option is out of money but the call option is in the money. To convert the in-the-money call option into the out-of-the-money put option, we use the put-call parity,

$$e^{rT}C(K_0, K) = e^{rT}P(K_0, T) + F - K_0. \tag{17}$$

If we plug this equality into equation (1) to convert all option prices into out-of-money option prices, we have

$$\sigma^2 = \frac{2}{T}\sum \frac{\Delta K}{K_i^2} e^{rT} P(K_i, T) + \frac{\Delta K_0}{TK_0^2}(F - K_0) - \frac{1}{T}\left[\frac{F}{K_0} - 1\right]^2, \tag{18}$$

where the second term on the right hand side of equation (18) 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.

Furthermore, we can approximate the interval,

$$\Delta K_0 = \frac{K_{-1} - K_0}{2} \approx F - K_0, \tag{19}$$

where K_{-1} denotes the available strike price just above the forward level. We obtain the above approximation if we assume that the forward level is in the middle of the two adjacent strike prices. Then, the last two terms cancel out in equation (18) if we use all out-of-money options. Thus, the VIX definition matches the theoretical relation for the quadratic return variation up to the jump-induced error term ε, and of course, the error induced by discretization of strikes. Another way to reconcile the adjustment term is to expand $\ln F_T$ around K_0 instead of around F for the derivation in the appendix.

Therefore, the new VIX index has very concrete economic meanings. It can either be regarded as the price of a portfolio of options, or regarded as an accurate approximation of the variance swap rate up to the discretzation error and the error induced by jumps.

2.3. Practical motivation for the switch

CBOE's switch from the old VIX (VXO) to the new VIX is motivated by a series of theoretical and also practical concerns. First, until very recently, the exact economic meaning of the VXO, or the at-the-money implied volatility, is not clear. In contrast, the definition of the new VIX is directly linked to the price of a portfolio of options. The economic meaning of the new VIX is much more concrete. Second, although we now understand that the VXO approximates accurately the volatility swap rate, it remains true that the volatility swap contracts are very difficult to replicate, much more difficult than the replication of the variance swap contract. Therefore, despite its popularity as a general volatility reference index, so far no derivative products have been launched on this VXO index. This is quite unique among indexes because almost all popular indexes have derivative products launched on them. In contrast, just a few months after the CBOE switched the VIX definition, they started planning to launch futures and options contracts on the new VIX. The currently planned launching date is March 28, 2004. Actually, forward contracts on the VIX (squared) are already available over the counter, offered by major investment banks.

3. Historical Behaviors

Based on historical data on daily closing option prices on S&P 500 index and S&P 100 index, the CBOE has back-calculated the VIX up to 1990 and the VXO up to 1987. We choose the overlapping sample period from January 2, 1990 to September 15, 2003 (5,005 daily observations) and analyze the historical behaviors of two indexes. We also download the returns on the S&P 100 index (OEX) and the S&P 500 index (SPX) from January 2, 1990 to October 15, 2003. The extra month for the index returns are used to compute the ex post realized return variance. We analyze how the volatility indexes correlate with the index returns and index return realized volatilities.

3.1. The time series of the two volatility indexes and realized variance

The top two panels in Figure 1 plot the time series of the new VIX and old VXO and their differences. The top left panel plots the two volatility index VIX and VXO. We see that they overlap one another well, indicating that they capture similar sentiments of the stock market. The top right panel plots the difference between VIX and VXO, which is close to zero in the early years and becomes mostly negative in more recent times. The Jensen inequality would imply a positive difference if the underlying security is the same. The observed negative difference implies that the S&P 100 index returns are more volatile than the S&P 500 index returns.

We also compute the ex post realized variance of the S&P 500 index returns and the S&P 100 index returns at each date over the subsequent 30 days. The bottom panels in Figure 1 plot the time series of the realized return variance (in annualized volatility percentage) and their differences. The realized return variance is more noisy, but the relative magnitude show similar patterns. Indeed, the realized variance on S&P 500 index returns is lower than the realized variance on S&P 100 index returns during the recent years.

Table 1 reports the summary statistics on the volatility indexes and the realized return variance. The two volatility indexes VIX and VXO differ by about one percentage point. The average difference between the realized return variance on S&P 500 and 100 indexes is slightly less than one percentage point. Therefore, the level difference between VIX and VXO can mostly be traced back to their underlying's volatilities.

Comparing the volatility index with the realized variance, we find that on average, the volatility index constructed from the options market is five percent higher than the realized volatility. To test the statistical significance of the difference between the volatility index and the realized volatility, we construct the following t-statistic,

$$t\text{-}\mathtt{stat} = \sqrt{N}\frac{\overline{X}}{S_X}, \tag{20}$$

where $N = 5,005$ denotes the number of observation, 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 the S&P 500 index return at 13 and for the S&P 100 index return as 13.8, both of which are strongly statistically significant.

The volatility levels exhibit moderate positive skewness and extra kurtosis, but the extra kurtosis for the daily differences is much larger. When we take logs on the volatility, the non-normalities for both the log levels and the log differences decline dramatically.

Figure 2 plots the cross-correlations between the index returns and the daily changes in the two volatility indexes. We find a strong negative instantaneous correlation between the two, but not significantly different from zero at other leads and lags. The same pattern holds for both the correlation between VIX and the S&P 500 index return and the correlation between VXO and the S&P 100 index return.

3.2. Weekday effects

It is well-known that volatilities during business days are on average higher than volatilities during market closes. The VIX index is interpolated to be the risk-neutral return variance for the next 30 actual days. However, different starting dates generate different number of business days for the subsequent 30-day period. In particular, starting on Mondays through Thursdays includes eight weekend (Saturdays and Sundays) days when the market is closed, but starting on Friday includes on extra Saturday. Thus, we expect the average VIX level on Fridays to be lower than the average VIX level on other weekdays.

To test this hypothesis, we sort the data based on weekdays and take the sample average within each weekday. The left panel in Figure 3 plots these sample averages for the five weekdays. Consistent with our hypothesis, the average VIX level on Fridays is lower than the average VIX on any other weekdays. To check its statistically significance, we compute the Newey-West serially adjusted standard deviation for the difference between the averages on other weekdays and the average on Fridays and compute the t-statistics for the difference according to equation (20). The t-statistics range from 10 to 16, highly significant.

Interestingly, the average VIX level also declines monotonically, albeit by a smaller magnitude, from Monday to Thursday. This decline cannot be explained by the weekend story.

Although VIX is computed based on actually number of days, the old VXO is computed using some business-day adjustment. As a result, we do not observe any obvious weekday pattern for VXO. The bar chart on the right panel of Figure 3 plots the sample averages of VXO on each weekday.

3.3. The FOMC meeting day effect

It has been found that Treasury bond and bill volume, bid-ask spreads, and volatility increase dramatically around FOMC meeting dates. We investigate whether the two volatility indexes show any apparent changes around the FOMC meeting days. For this purpose, we download the FOMC meeting day logs from Bloomberg. During our sample period, there are altogether 110 scheduled FOMC meetings, about ten meetings per year. With a list of these meeting days, we sort the two volatility indexes around the FOMC meeting days and compute the average index level ten days before and ten days after the FOMC meeting days. Figure 4 plots the sample averages of VIX (left panel) and VXO (right panel) around the FOMC meeting days.

For both VIX and VXO, we observe that the volatility level drops markedly after the FOMC meeting day. For VIX, the volatility reaches the highest level the day before the meeting and drops to the lowest level three days after the meeting. For VXO, the vol reaches the highest level four days before the meeting and drops to the bottom four days after the meeting.

To investigate the significance of the drop, we measure the difference between the vol index one day before and one day after the meeting. The mean difference is 0.7625 for VIX and 0.4745 for VXO, both in percentage volatility points. The t-statistics for the two differences are 4.77 and 2.70, respectively.

Before FOMC meeting day, people disagree on wether the Fed will change the federal funds rate target. While the degree of uncertainty may not vary much during the last few days, the uncertainty in annualized volatility terms increases. The fact that the option-implied equity volatility increases implies that the uncertainty on interest rates has a definite impact on the volatility in the equity market.

This uncertainty is resolved right after the meeting and hence the vol index drops rapidly after the meeting.

Since we can consider a variance swap contract on S&P 500 index and regard the VIX as the variance swap rate of this contract, we investigate whether timing the variance swap investment around the FOMC meeting days makes a difference. Figure 5 plots the average ex post payoff for longing the swap contract around the FOMC meeting days. The payoff is defined as the difference between the ex post realized variance and the VIX squared: $(RV_{t,T} - VIX_t^2)$, where RV denotes the annualized realized variance during the subsequent 30-day period. We find that the average payoffs are negative by longing the swap on any days. That is, shorting the swap contract generates positive payoffs on average. In particular, we find that shorting the swap contract four pays prior to the FOMC meeting days generates the highest payoff, and that shorting the variance swap four days after the FOMC meeting days generates the lowest payoff. The difference in average payoffs between investments in these two days is statistically significant, with a t-statistics of 7.74.

3.4. The expectation hypothesis and the market index return variance risk premium

The VIX level reflects the risk-neutral expectation of the quadratic variation of the index return.

$$VIX_t^2 = \frac{1}{T-t}\mathbb{E}_t^{\mathbb{Q}}\left[\ln F_T/F_t, \ln F_T/F_t\right] = \mathbb{E}_t^{\mathbb{P}}\left[M_{t,T}\, RV_{t,T}\right]. \tag{21}$$

where $T-t$ is 30 days and $M_{t,T}$ denotes the state price density that changes the measure from $\mathbb{P}$ to $\mathbb{Q}$.

Rearrange equation (21), we have

$$\begin{aligned} 1 &= \mathbb{E}_t^{\mathbb{P}}\left[M_{t,T}\frac{RV_{t,T}}{VIX_t^2}\right] = \mathbb{E}_t^{\mathbb{P}}\left[M_{t,T}\right]\mathbb{E}_t^{\mathbb{P}}\left[\frac{RV_{t,T}}{VIX_t^2}\right] + Cov_t\left[M_{t,T}, \frac{RV_{t,T}}{VIX_t^2}\right] \\ &= \mathbb{E}_t^{\mathbb{P}}\left[\frac{RV_{t,T}}{VIX_t^2}\right] + Cov_t\left[M_{t,T}, \frac{RV_{t,T}}{VIX_t^2}\right]. \end{aligned} \tag{22}$$

Therefore, if we measure the ratio of the realized variance to the VIX square and compute the average value, any significant deviation of this average value from unity reflects the covariance between the pricing kernel and this ratio. The negative of this covariance measures the premium that the market charges on the index return variance risk.

We compute the average of the ratio (RV/VIX^2) at 0.6027. The t-statistic against the null hypothesis of the ratio being one is -12.38. Thus, the ratio is significantly less than one. The variance risk premium on the S&P 500 index is significant negative.

If we consider the VIX in the variance swap context and regard $RV_{t,T}$ as future payoff and VIX as the forward price, then the ratio defines the raw excess return. The negative risk premium implies that being long on the variance swap generates a significant negative excess return (or raw excess return less than one).

Analogously, we can regard VXO as the volatility swap rate, which is the risk-neutral expected value of the OEX return volatility. A similar relation to equation (22) exists between the VXO level and the realized OEX return volatility. We estimate the average ratio $(VXO/\sqrt{RV})$ at 0.7484. The t-statistics against the null value of one is -15.51, again strongly significant and showing that the volatility or variance risk premia on both SPX and OEX are strongly negative.

To test whether the variance risk premium is time varying, we run the following expectation-hypothesis (EH) regressions:

$$V_{t,T} = a + bVIX_t + e_{t,T}, \tag{23}$$

where we choose to represent the volatility in three different forms: volatility, variance, and log volatility. Due to the overlapping data, if we use daily data for the regression, the R-square of the regression will be artificially high. To correct this problem, we run the regression on data sampled 30 days apart. To make full use of the data, we run the regression at 30 different starting dates and report the sample averages of the parameter estimates, standard deviations, and R-squares. We find that the regression results differ very little when we change the starting dates so that the sample averages do not differ much for each single run.

Table 2 report the regression results. We find that for both indexes, the R-square is highest when the regression is on log volatility, lowest when the regression is on variance. The slope estimate is also closest to one when the regression is on log volatility. Indeed, the slope estimate is not significantly different from one when the regression is run on log volatility, but significantly lower than one when the regression is on volatility or variance. Therefore, if we formulate the regression in logarithm, we cannot reject the null value of one for the slope and the null hypothesis that the risk premium in log terms is a constant or independent series.

3.5. The information content of the volatility indexes

In this section, we estimate GARCH(1,1) processes on the S&P 500 index return and the S&P 100 index returns during the common sample period from 1990 to 2003. Then, we compare the relative information content of the GACRH volatility and the option-volatility index. Table 3 reports the regression results.

First, we regress the realized log volatility solely on the GARCH log volatility and compare the R-squares with the regressions on log option-implied volatility index. Given our previous findings on the high R-squares when using logarithms of volatility, we only run regressions in log forms in this section, using the same method as before. The regression results show that the volatility index has higher forecasting power in terms of the R-square than the GARCH volatility. When we use both the volatility index and the GARCH volatility as explanatory variables, the R-square does not improve much over the R-square using the volatility index alone. Similar results hold for both S&P 500 and 100 index return variance. Therefore, we conclude that the VIX and VXO are quite efficient forecasts on the future realized variance. The GARCH volatility does not seem to provide much extra information in addition to that in the volatility index. Furthermore, the fact that the slope estimate is close to one when running the regression on log terms implies that variation rates for the the realized variance and for the option-implied volatility index have an approximate one-to-one correspondence.

Comparing the two indexes, we find that that the regressions on VXO generate slightly higher R-squares than the corresponding regressions on VIX. This difference can come either from the difference in the underlying index (OEX versus SPX), or from the definition of the volatility index (at-the-money vol versus variance swap rate). Our experience indicates that at-the-money implied volatility provides most of the information on the realized volatility. Incorporating the out-of-the-money options in VIX can reduce the average bias in terms of approximating the return variance, but given the less liquidity for out-of-the-money options, the return variance swap rate formed with both at-the-money and out-of-the-money options tend to have slightly reduced forecasting power than using the at-the-money implied volatility alone.

Figure 6 plots the cross-correlation between the two volatility indexes and their respective realized volatilities. Due to the overlapping sample and the fact that we measure the correlation on volatility levels, we observe high cross-correlation. The most illustrative is the peak that happens at the 30-day lags for both indexes. This peak implies that the current VIX and VXO levels are the most correlated with the realized variance during the past 30-days. Measuring the cross-correlation on changes generate much smaller and also much noisier correlation estimates. However, the highest correlation estimates remain at the 30-day lag for the realized variance.

Figure 7 overlaps the volatility index (solid lines) with the realized volatility (dashed lines). We observe that almost all large spikes in the volatility indexes are precedented by large spikes in the realized volatility. Both Figure 6 and Figure 7 show that the realized variance forecast the volatility index better than the reverse and that many times the investors mend the fence *after* the sheep have fled.

3.6. The excess return on entering a variance-swap contract

Since VIX squared can be regarded as the variance swap rate on SPX, if we long a variance swap contract today, at maturity (30 days later) we will receive the realized variance and pay the fixed VIX-squared determined 30 days ago. For such a zero-financing investment, we compute the excess return as

$$ER_{t,T} = (RV_{t,T} - VIX_t^2)/VIX_t^2. \tag{24}$$

Analogously, we can regard VXO as the volatility swap rate on OEX. if we long a volatility swap contract, we can compute the excess return as

$$ER_{t,T} = (\sqrt{RV}_{t,T} - VXO_t)/VXO_t. \tag{25}$$

Alternatively, we can represent the excess return in continuous compounding format,

$$LER_t = \ln(1 + ER_t). \tag{26}$$

Table 4 reports the summary statistics of the excess returns on the variance swap and volatility swap, respectively. Longing both swap contracts generate negative excess returns on average. For the variance swap contract on SPX, the average return is -39.7 percent per month in simple compounded form and -65.9 percent per month in continuously compounded form. The average return on longing the OEX volatility swap contract is also negative at -25.2 percent per month in simple compounding term and -32.4 percent per month in continuous compounding term.

The left panels in Figure 8 plots the histogram of the simple compounded excess returns on the swap contracts. The excess returns are predominantly negative for both contracts, but there are some large positive return realizations so that the histogram is highly skewed to the right. In the right panels, we plot the histogram of the log excess returns, which are much closer to a normal distribution shape. The skewness and kurtosis estimates reported in Table 4 tell a similar story. Although the simple compounding excess returns generate large and positive skewness and kurtosis estimates. These estimates for the log return are close to zero.

Figure 9 plots the time series of the excess return ER_t and the log excess return LER_t on the SPX variance swap contracts and OEX volatility swap contracts. The excess returns are predominantly negative, but we observe more positive returns during the later years.

Given the large and negative mean excess returns, it seems quote profitable to short the two swap contracts. The last column in Table 4 reports the annualized Sharpe ratio for shorting the two swap contracts. The Sharpe ratios are between three and four, indeed very large. We compute the Sharpe ratio using 30-day apart non-overlapping data. We report the average of the estimates from different starting days. In computing the Sharpe ratio, the standard deviation for the non-overlapping returns are adjusted for serial dependence according to Newey and West (1987).

The high average profitability from shorting the variance and volatility swap contracts imply that investors are willing to receive a large average negative return by longing the variance swap contract. Why so? We investigate whether the market portfolio risk, or beta risk, explains this large negative risk premium. For this purpose, we run the following capital asset pricing model (CAPM) regression,

$$ER_t = \alpha + \beta\,(R_t^m - R_f) + e_t, \tag{27}$$

where $(R_t^m - R_f)$ denotes the excess return to the market portfolio. If the CAPM theory holds, we would expect to obtain a highly negative beta estimate for the long variance or volatility swap return. This is actually possible given the well-documented correlation between market index returns and the return variance. Nevertheless, if the CAPM can fully account for the risk premium, we would expect the estimate for the intercept α not significantly different from zero. The intercept α represents the excess return to a market-neutral strategy that involves long a variance swap contract and short beta of the market portfolio.

We use log returns on both sides. 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 this excess return is publicly available at Kenneth French's data library on the web from 1926 to 2003.[2] We match the sample period with our data and run the regression on monthly returns over non-overlapping data using the generalized methods of moments.

The regression estimates are as follows, with standard errors of the estimates in the parentheses below the estimates,

$$
\begin{array}{lcccccll}
SPX: & LER_t & = & -0.6165 & - & 4.6991 & (R_t^m - Rf) & +e, \quad R^2 = 14.05\%. \\
 & & & (0.0088) & & (0.1607) & & \\
OEX: & LER_t & = & -0.3047 & - & 2.3875 & (R_t^m - Rf) & +e, \quad R^2 = 15.49\%. \\
 & & & (0.0036) & & (0.0763) & &
\end{array}
\tag{28}
$$

First, we observe that the beta estimates for both swap contracts are highly negative, consistent with the general observation that the index return and volatility are negatively correlated. However, this negative beta cannot fully explain the negative risk premium on the variance and volatility risks. The estimates for the intercepts, or the market-neutral excess returns, remain strongly negative. Actually, the magnitudes of α are not much smaller than the raw excess returns (-65.9 and -32.4 percent for SPX and OEX, respectively). Thus, we conclude that the CAPM only gets the sign right, but cannot fully account for the large negative risk premium on index return variance risk.

[2]The web address is: `http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html`.

It is well-known that the return volatility shows strong mean reverting property and hence is predictable. Now that we have variance and volatility contracts, investors can exploit such predictability and directly convert them into dollar returns. We investigate whether the excess returns on SPX variance swap contract and the OEX volatility contract are predictable.

First, we look at the monthly autocorrelation estimates for the excess returns, which are reported under “Auto” in Table 4. These estimates are estimated using non-overlapping 30-day part data. The autocorrelation is about 0.14 for the excess returns on the SPX variance swap and 0.07 for the excess returns on the OEX volatility swap. Both numbers are small. When we run AR(1) regressions on the excess returns, we obtain an average R-square estimate around two percent for returns on SPX variance swap and below one percent for returns on OEX volatility swap. Thus, the forecasting power on the excess return is very low. Although the volatility level is strongly predictable, investors have priced this predictability into the variance or volatility swap contract so that the excess returns on these swap contracts are not strongly predictable.

Second, we investigate whether we can forecast the excess return on the variance and volatility swap contract using the index returns. Figure 10 plots the cross-correlation between the excess return to the variance or volatility swap contract and the monthly return on the underlying stock index, based on monthly sampled and hence non-overlapping data. We find that the stock index return and the return on the swap contracts have strong negative contemporaneous correlations, but the non-overlapping series do not seem to have any lead-lag effects.

4. Conclusion

The new VIX differs from the old VXO in two key areas: They use different underlyings (SPX for the new VIX versus OEX for the old VXO), and they use different formulae in extracting volatility information from the options market. The new VIX definition represents a model-free approximation of the return variance swap rate, whereas the old VXO approximates the volatility swap rate under certain assumptions. The CBOE decides to switch from VXO to VIX mainly because the new VIX has more concrete economic meanings and the variance swap the new VIX represents has a straightforward replicating portfolio in options. In contrast, replicating the volatility swap contract the VXO represents is much more difficult.

The historical behavior of the two volatility indexes are very similar. On average, the VXO is about one percentage point higher the VIX. We observe about the same magnitude of difference between the realized monthly volatility of S&P 100 index and S&P 500 index. We also find a strong negative contemporaneous correlation between index returns and changes in the volatility indexes.

Since the VIV is based on a maturity of 30 calendar days, we find that on average the VIX level on Fridays is about half a percentage lower than the VIX level on Mondays. Furthermore, both VIX and VXO drop about one percentage point after the FOMC meeting announcements.

When comparing the volatility index with the realized return variance, we find that on average the index is about five percentage point higher than the realized volatility. Nevertheless, when we regress the log realized volatility on the log option-implied volatility index, the regression slope estimates are not significantly different from one for both VIX and VXO, the null value if we assume that the log premium on the volatility risk is constant or independent of the log variance swap rate.

We also investigate the relative information content of the two volatility indexes versus GARCH type volatility estimates. We find that the volatility indexes are better forecasts of the realized volatility than GARCH forecasts. Furthermore, once the volatility index is included, GARCH volatility does not add additional information to the volatility forecasts.

We observe that shorting the SPX variance swap contracts or OEX volatility swap contracts, as represented by VIX and VXO, respectively, generates highly positive returns and high Sharpe ratios. Furthermore, the beta risk only gets the sign right, but cannot fully explain the large magnitude of the negative variance risk premium on the two stock indexes.

Appendix A. Proof

Let $f(F_t)$ denote a general payoff function of F_T. By Itô's lemma, we have

$$
\begin{aligned}
f(F_T) &= f(F_0)+\int_0^T f'(F_{t-})dF_t+\frac{1}{2}\int_0^T f''(F_{t-})d\sigma_{t-}^2 \\
&\quad +\int_0^T\int_{\mathbb{R}}^0 [f(F_{t-}e^x)-f(F_{t-})-f'(F_{t-})F_{t-}(e^x-1)]\mu(dx,dt), \qquad \text{(A1)}
\end{aligned}
$$

Apply (A1) to the function $f(F)=\ln F$, we have

$$
\ln(F_T) = \ln(F_0)+\int_0^T \frac{1}{F_{t-}}dF_t-\frac{1}{2}\int_0^T \sigma_{t-}^2 dt+\int_0^T\int_{\mathbb{R}}^0 [x-e^x+1]\mu(dx,dt).
$$

Add and subtract $2[\frac{F_T}{F_0}-1]+\int_0^T x^2\mu(dx,dt)$ and rearrange, we obtain a representation for the quadratic variation for the asset return,

$$
\begin{aligned}
\int_0^T \sigma_{t-}^2 dt+\int_0^T x^2\mu(dx,dt) &= 2\left[\frac{F_T}{F_0}-1-\ln\left(\frac{F_T}{F_0}\right)\right]+2\int_0^T\left[\frac{1}{F_{t-}}-\frac{1}{F_0}\right]dF_t \\
&\quad -2\int_0^T\int_{\mathbb{R}}^0\left[e^x-1-x-\frac{x^2}{2}\right]\mu(dx,dt). \qquad \text{(A2)}
\end{aligned}
$$

A Taylor expansion with remainder of $\ln F_T$ about the point F_0 implies,

$$
\ln F_T=\ln F_0+\frac{1}{F_0}(F_T-F_0)-\int_0^{F_0}\frac{1}{K^2}(K-F_T)^+dK-\int_{F_0}^{\infty}\frac{1}{K^2}(F_T-K)^+dK. \qquad \text{(A3)}
$$

Plug (13) into the left hand side of (A2) and plus (A3) into the right hand side of (A2), we have

$$
\begin{aligned}
[\ln F_T,\ln F_T] &= 2\left[\int_0^{F_0}\frac{1}{K^2}(K-F_T)^+dK+\int_{F_0}^{\infty}\frac{1}{K^2}(F_T-K)^+dK\right] \\
&\quad +2\int_0^T\left[\frac{1}{F_{t-}}-\frac{1}{F_0}\right]dF_t \\
&\quad +2\int_0^T\int_{\mathbb{R}}^0\left[e^x-1-x-\frac{x^2}{2}\right]\mu(dx,dt). \qquad \text{(A4)}
\end{aligned}
$$

Thus, we can replicate the quadratic return variation of the asset return by (i) the payoff from a continuum of European out-of-the-money options: call options when $F>K$ and put options when $K\leq F$ (first line), (ii) the payoff from a dynamic trading strategy (second line) and (iii) a higher order error term generated from the discontinuity of the futures price dynamics (third line).

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

$$\mathbb{E}_0^Q[\ln F_T, \ln F_T] \quad = \quad e^{rT}\int_0^\infty \frac{2P_0(K,T)}{K^2}dK - 2\mathbb{E}_0^Q\int_0^T\int_{\mathbb{R}}^0\left[e^x - 1 - x - \frac{x^2}{2}\right]\nu_t(x)dxdt,$$

where $P_0(K,T)$ denote the time-0 value of the European out-of-the-money option at strike price K and expiry T. By the martingale property of the futures prices, the expected value of the payoff from the the dynamic futures trading is zero under the risk-neutral measure.

In equation (A3), we expand $\ln F_T$ around the current forward level F_0 to obtain the expected value of the return quadratic variation as a portfolio of out-of-money options, up to an error term. Alternatively, we can expand $\ln F_T$ around $K_0 \leq F_0$,

$$\ln F_T = \ln K_0 + \frac{1}{K_0}(F_T - K_0) - \int_0^{K_0}\frac{1}{K^2}(K - F_T)^+ dK - \int_{K_0}^\infty \frac{1}{K^2}(F_T - K)^+ dK. \tag{A5}$$

If we plug this expansion into the quadratic variation equation in (A2), we would have generate a forth term,

$$\begin{aligned}
[\ln F_T, \ln F_T] \quad = \quad & 2\left[\int_0^{K_0}\frac{1}{K^2}(K - F_T)^+ dK + \int_{K_0}^\infty \frac{1}{K^2}(F_T - K)^+ dK\right] \\
& + 2\int_0^T\left[\frac{1}{F_{t-}} - \frac{1}{F_0}\right]dF_t \\
& + 2\int_0^T\int_{\mathbb{R}}^0\left[e^x - 1 - x - \frac{x^2}{2}\right]\mu(dx,dt) \\
& + 2\left[\frac{F_T}{F_0} - \frac{F_T}{K_0} + \ln\frac{F_0}{K_0}\right].
\end{aligned} \tag{A6}$$

Taking expectations on this extra term, we have

$$\begin{aligned}
Extra \quad &= \quad 2\mathbb{E}_0^{\mathbb{Q}}2\left[\frac{F_T}{F_0} - \frac{F_T}{K_0} + \ln\frac{F_0}{K_0}\right] \\
&= \quad 1 - \frac{F_0}{K_0} + \ln\frac{F_0}{K_0}.
\end{aligned} \tag{A7}$$

Taylor expand the log term, we have

$$\ln\frac{F_0}{K_0} \approx \frac{F_0 - S_0}{K_0} - \frac{1}{2}\frac{(F_0 - K_0)^2}{K_0^2}, \tag{A8}$$

Plug (A8) into (A7), we have

$$Extra \approx -\frac{1}{2}\frac{(F_0 - K_0)^2}{K_0^2}. \tag{A9}$$

which matches the extra term in the VIX definition in using $K_0 \leq F_0$ instead of F_0 as the benchmark point.

References

Andrews, Donald, 1991, Heteroskedasticity and autocorrelation consistent covariance matrix estimation, *Econometrica* 59, 817–858.

Black, Fisher, and Myron Scholes, 1973, The pricing of options and corporate liabilities, *Journal of Political Economy* 81, 637–654.

Blair, Bevan J., Ser-Huang Poon, and Stephen J. Taylor, 2000, Forecasting s&p 100 volatility: The incremental information content of implied volatilities and high frequency index returns, Working paper, Lancaster University.

Brenner, Michael, and Marti Subrahmanyam, 1988, A simple formula to compute the implied standard deviation, *Financial Analysts Journal* 44, 80–83.

Canina, Linda, and Stephen Figlewski, 1993, The information content of implied volatility, *Review of Financial Studies* 6, 659–681.

Carr, Peter, and Roger Lee, 2003, At-the-money implied as a robust approximation of the volatility swap rate, Working paper, Bloomberg LP.

Carr, Peter, and Dilip Madan, 1998, Towards a theory of volatility trading, in Robert Jarrow, eds.: *Risk Book on Volatility* (Risk, New York).

Carr, Peter, and Liuren Wu, 2003, Predictability of return variance swap: An asset allocation perspective, Working paper, Fordham University.

Chiras, D., and S. Manaster, 1978, The information content of option prices and a test of market efficiency, *Journal of Financial Economics* 6, 213–234.

Christensen, Bent Jesper, and Charlotte Strunk Hansen, 2002, New evidence on the implied-realized volatility relation, *European Journal of Finance* 8, 187–205.

Christensen, B. J., and N. R. Prabhala, 1998, The relation between implied and realized volatility, *Journal of Financial Economics* 50, 125–150.

Day, Theodore E., and Craig M. Lewis, 1988, The behavior of the volatility implicit in option prices, *Journal of Financial Economics* 22, 103–122.

Day, Theodore E., and Craig M. Lewis, 1992, Stock market volatility and the information content of stock index options, *Journal of Econometrics* 52, 267–287.

Day, Theodore E., and Craig M. Lewis, 1994, Forecasting futures market volatility, *Journal of Derivatives* 1, 33–50.

Demeterfi, Kresimir, Emanuel Derman, Michael Kamal, and Joseph Zou, 1999, A guide to volatility and variance swaps, *Journal of Derivatives* 6, 9–32.

Fleming, Jeff, 1998, The quality of market volatility forecast implied by s&p 500 index option prices, *Journal of Empirical Finance* 5, 317–345.

Gwilym, Owain Ap, and Mike Buckle, 1999, Volatility forecasting in the framework of the option expiry cycle, *European Journal of Finance* 5, 73–94.

Hansen, Charlotte Strunk, 2001, The relation between implied and realized volatility in the Danish option and equity markets, *Accounting and Finance* 41, 197–228.

Hol, Eugenie, and Siem Jan Koopman, 2000, Forecasting the variability of stock index returns with stochastic volatility models and implied volatility, Working paper, Free University Amsterdam.

Hull, John, and Alan White, 1987, The pricing of options on assets with stochastic volatilities, *Journal of Finance* 42, 281–300.

Jacod, Jean, and Albert N. Shiryaev, 1987, *Limit Theorems for Stochastic Processes*. (Springer-Verlag Berlin).

Jorion, Philippe, 1995, Predicting volatility in the foreign exchange market, *Journal of Finance* 50, 507–528.

Lamoureux, Christopher G., and William D. Lastrapes, 1993, Forecasting stock-return variance: Toward an understanding of stochastic implied volatilities, *Review of Financial Studies* 6, 293–326.

Latane, Henry A., and Richard J. Rendleman, 1976, Standard deviation of stock price ratios implied in option prices, *Journal of Finance* 31, 369–381.

Neely, Christopher J., 2003, Forecasting foreign exchange rate volatility: Is implied volatility the best we can do?, Working paper, Federal Reserve Bank of St. Louis.

Newey, Whitney K., and Kenneth D. West, 1987, A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix, *Econometrica* 55, 703–708.

Prokhorov, Yurii Vasilevich, and Albert Nikolaevich Shiryaev, 1998, *Probability Theory III : Stochastic Calculus*. (Springer-Verlag Berlin).

Shu, Jinghong, and Jin E. Zhang, 2003, The relationship between implied and realized volatility of S&P500 index, *Wilmott Magazine* January, 83–91.

Tabak, Benjamin Miranda, Eui Jung Chang, and Sandro Canesso de Andrade, 2002, Forecasting exchange rate volatility, Working paper, Banco Central do Brasil.

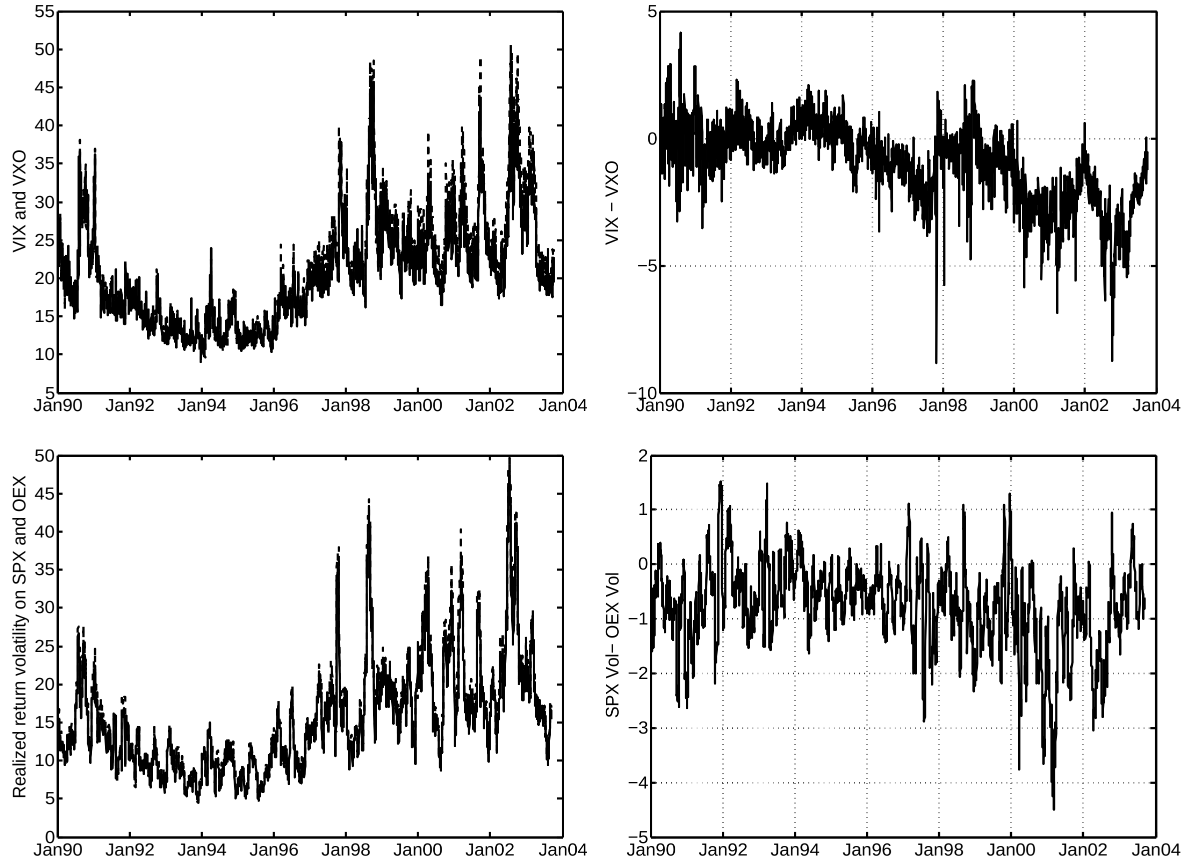

Figure 1. VIX, VXO, and realized return volatilities on S&P 500 and 100 indexes.

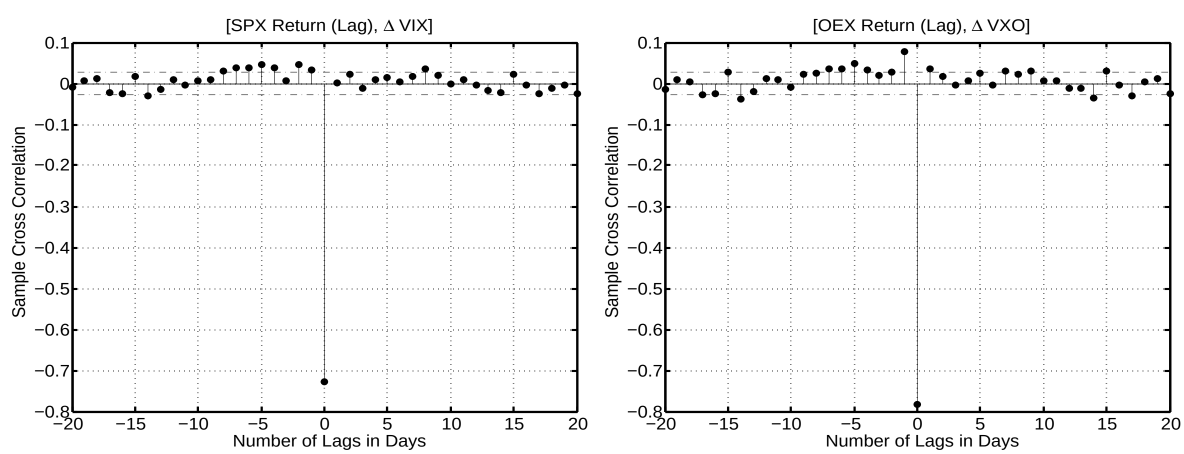

Figure 2. Cross-correlations between return and volatility.

The stem bars represent the cross-correlation estimates between the index returns and the corresponding volatility indexes. The two dashed lines in each panel denote the 95 confidence band.

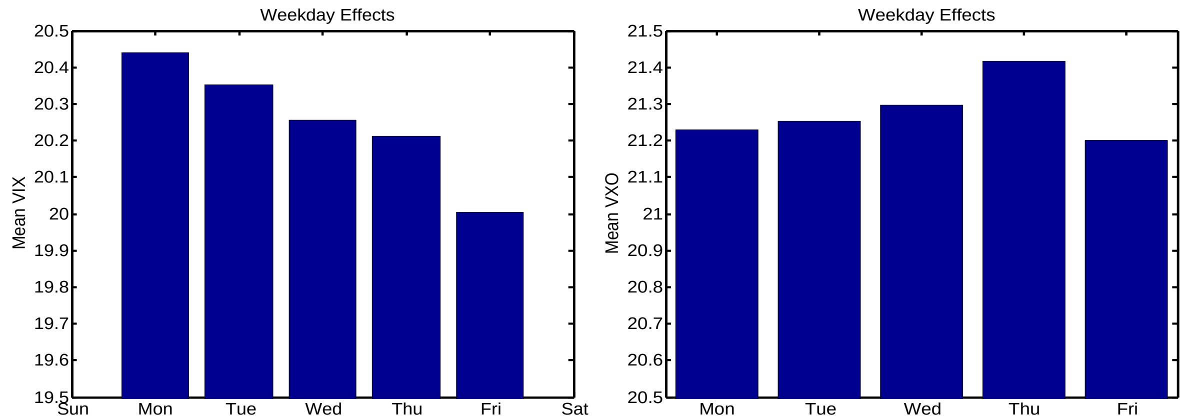

Figure 3. Weekday effects of VIX and VXO.
The bars plot the sample averages of VIX (left panel) and VXO (right panel) on each weekday.

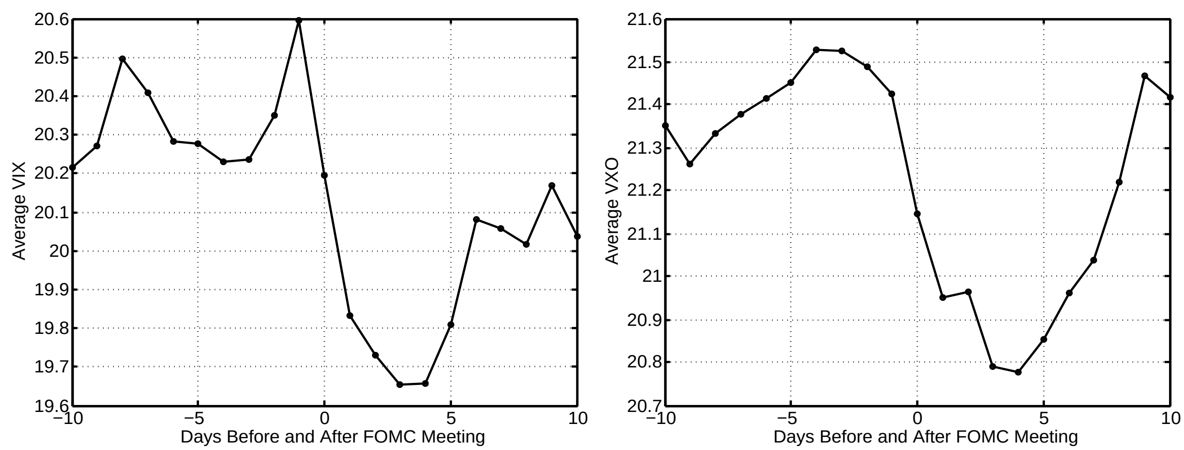

Figure 4. Volatility indexes around FOMC meeting days.
Lines represent the sample average of the VIX (left panel) and VXO (right panel) levels ten days before and ten days after the FOMC meeting days.

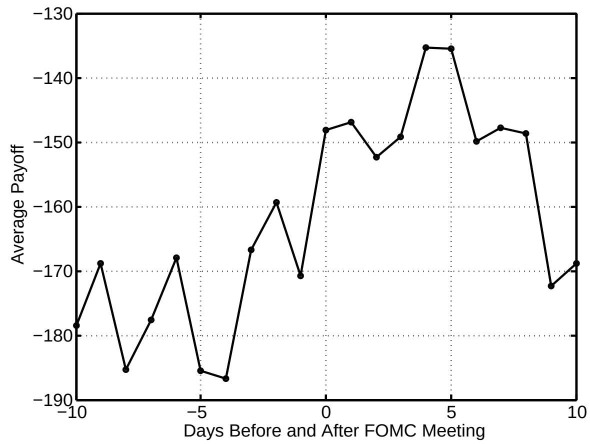

Figure 5. Payoffs to long variance swap contracts signed around FOMC meeting days.
The line represents the average payoff to long variance swap contracts ten days before and ten days after the FOMC meeting days.

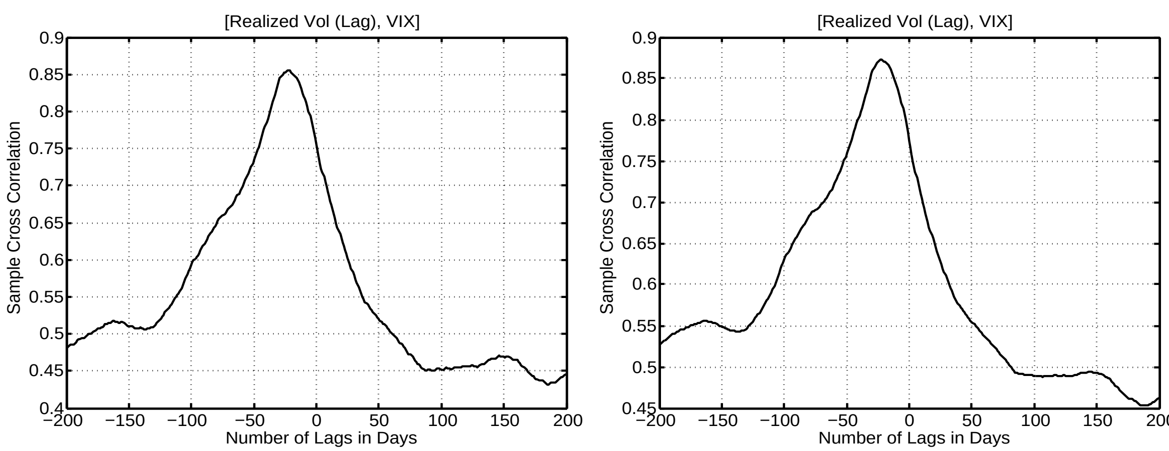

Figure 6. Cross-correlation between the volatility indexes with the realized volatilities.

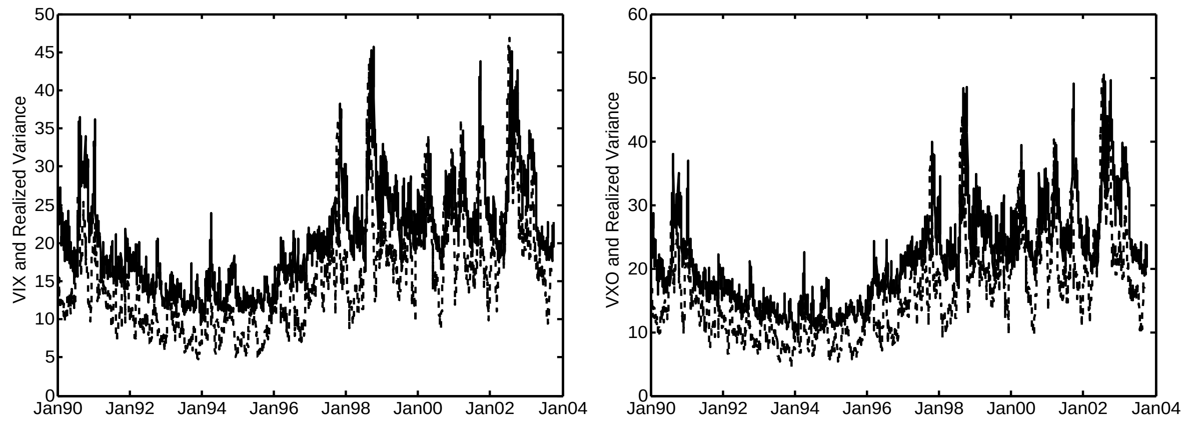

Figure 7. Comparing the volatility index and the realized volatility
The solid lines are the VIX (left) and VXO (right) from the options data. The dashed lines are the annualized realized 30-day volatility for the S&P 500 and S&P 100 index returns.

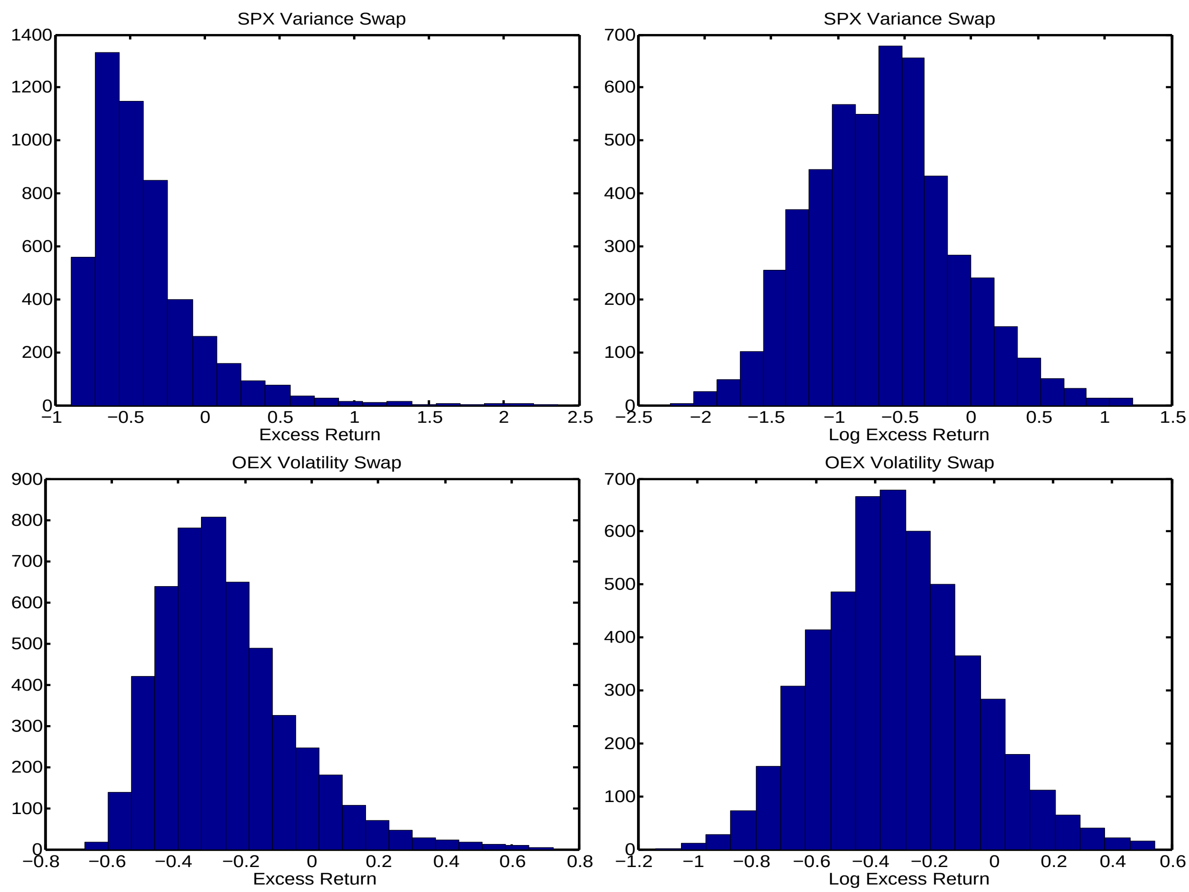

Figure 8. Histogram of the excess return on longing SPX variance swap and OEX volatility swap.

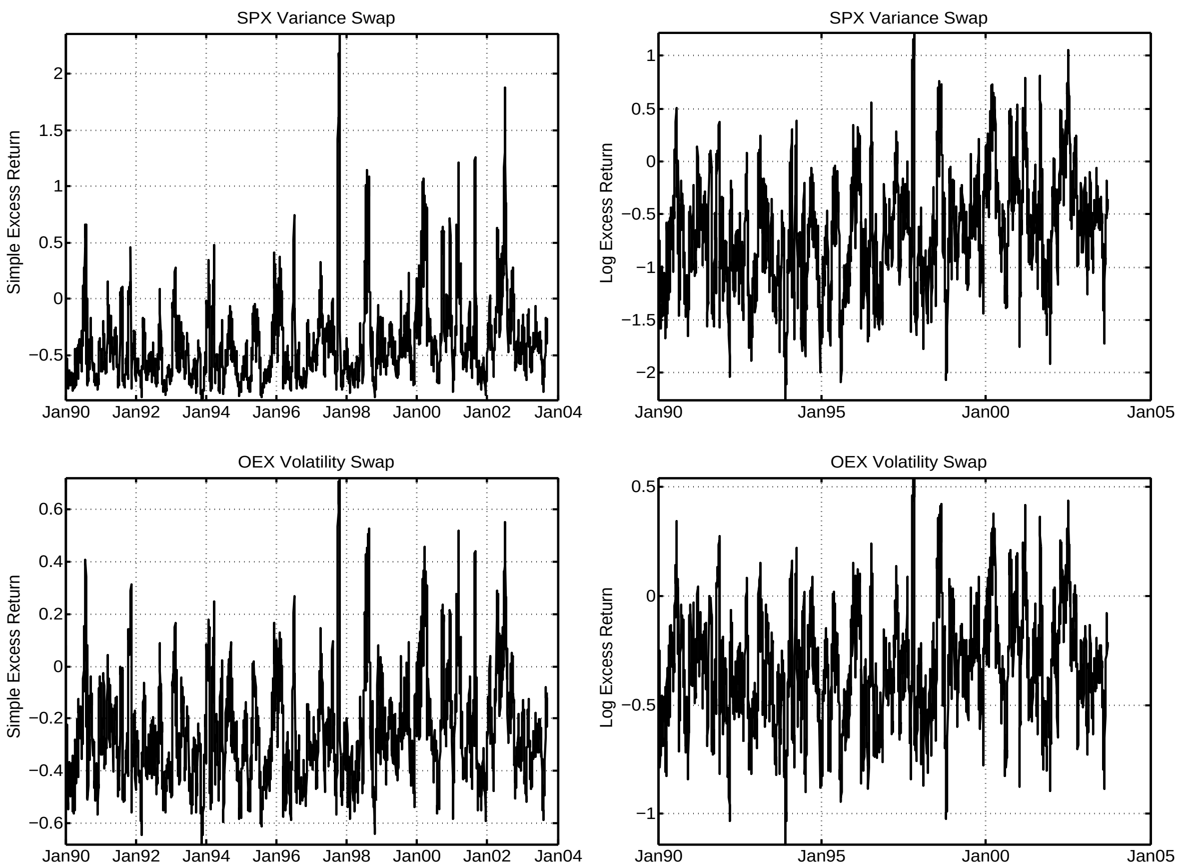

Figure 9. Time series of the excess returns on longing the SPX variance swap and OEX volatility swap.

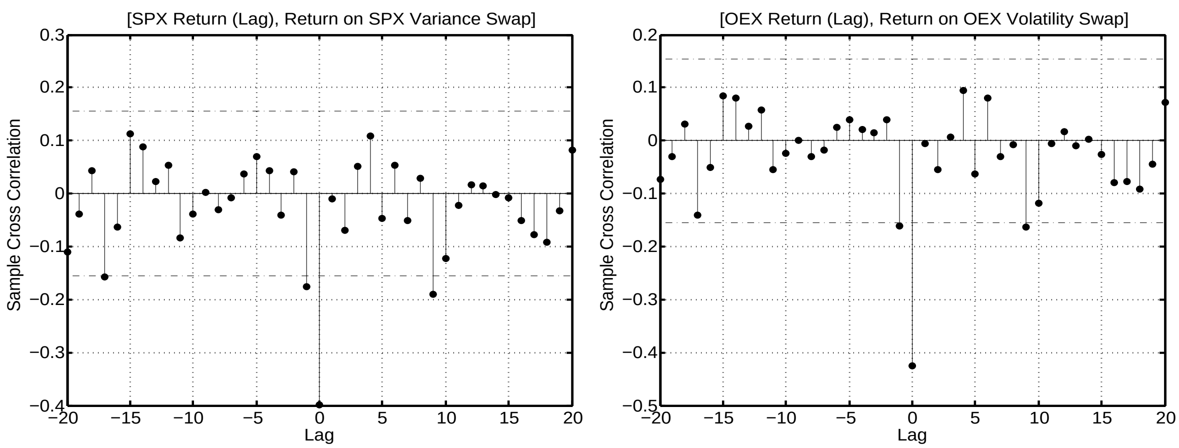

Figure 10. Cross-correlation between the stock index monthly returns and monthly returns on the variance/volatility swap.
The stem bars represent the cross-correlation estimates between the excess returns on the variance/volatility contracts and the corresponding stock index returns, based on monthly non-overlapping data. The two dashed lines in each panel denote the 95 confidence band.

Table 1
Summary statistics of volatility index and realized return variance

Entries are summary statistics on the daily series of the two volatility indexes, VIX and VXO, as well as the corresponding daily estimates of the ex post realized volatility. The realized volatility is annualized and is based on raw second moments (without demeaning). The sample has 5,005 daily observations from January 2, 1990 to September 15, 2003.

	VIX	SPX Vol	VXO	OEX Vol	VIX	SPX Vol	VXO	OEX Vol
	Levels				Daily Differences			
Mean	20.180	15.241	21.254	16.030	0.000	-0.000	0.001	-0.000
Stdev	6.486	7.087	7.391	7.527	1.060	0.863	1.216	0.907
Skewness	0.807	1.282	0.811	1.257	0.668	0.728	0.676	0.539
Kurtosis	0.519	2.061	0.538	1.865	9.638	30.288	12.916	27.775
Auto	0.987	0.993	0.986	0.993	-0.029	0.051	-0.089	0.060
	Log Levels				Daily Log Differences			
Mean	2.955	2.625	2.998	2.674	0.000	-0.000	0.000	-0.000
Stdev	0.314	0.442	0.342	0.446	0.047	0.055	0.049	0.055
Skewness	0.104	0.103	0.060	0.118	0.736	0.485	0.600	0.387
Kurtosis	-0.651	-0.383	-0.697	-0.424	6.729	17.840	7.157	18.549
Auto	0.989	0.992	0.990	0.992	-0.033	0.024	-0.077	0.027

Table 2
Expectation hypothesis on the variance risk premia

Series	Intercept		VIX/VXO		R-square
			A. S&P 500 index		
Variance	-12.756	(29.937)	0.658	(0.055)	0.456
Vol	-1.442	(1.179)	0.827	(0.056)	0.569
Log Vol	-0.685	(0.197)	1.120	(0.066)	0.631
			B. S&P 100 index		
Variance	1.415	(30.964)	0.617	(0.050)	0.480
Vol	-0.769	(1.124)	0.790	(0.050)	0.599
Log Vol	-0.519	(0.177)	1.065	(0.059)	0.663

Entries report the results of regressing the realized variance against the volatility index. Since we have daily estimates of monthly volatility, we sample the data monthly to avoid overlapping data. To make full use of the information, we run the regression with 30 different starting dates and report the averages of 30 estimates on the parameters, standard errors (in parentheses), and R-squares. For each index, we run the regression based on the variance, the volatility, and the log of the volatility.

Table 3
Information content of the volatility indexes

Series	Intercept		VIX/VXO		GARCH		R-square
			A. S&P 500 index				
VIX	-0.685	(0.197)	1.120	(0.066)	—	—	0.631
GARCH	0.180	(0.166)	—	—	0.907	(0.061)	0.568
Joint	-0.580	(0.201)	0.846	(0.147)	0.262	(0.125)	0.641
			B. S&P 100 index				
VXO	-0.519	(0.177)	1.065	(0.059)	—	—	0.663
GARCH	0.173	(0.169)	—	—	0.911	(0.061)	0.572
Joint	-0.500	(0.179)	0.969	(0.143)	0.098	(0.132)	0.666

Entries report the results of regressing the realized variance against the volatility index and the GARCH(1,1) volatility forecasts. Since we have daily estimates of monthly volatility, we sample the data monthly to avoid overlapping data. To make full use of the information, we run the regression with 30 different starting dates and report the averages of 30 estimates on the parameters, standard errors (in parentheses), and R-squares. For each index, we run the regression based on the log of the volatility.

Table 4
The excess return on long variance (volatility) swap contracts

Series	Mean	Std	Skewness	Kurtosis	Auto	Sharpe
			A. Variance Swap Contract on SPX			
ER	-0.397	0.377	2.301	8.345	0.134	3.337
LER	-0.659	0.541	0.253	0.033	0.153	3.745
			B. Volatility Swap Contract on OEX			
ER	-0.252	0.205	1.116	1.864	0.077	4.103
LER	-0.324	0.259	0.290	0.045	0.067	4.204

Entries report the summary statistics of the excess return on longing a variance swap contract on SPX and longing a volatility swap contract on OEX. ER denotes simply compounded excess return, defined as $ER = (RV - VIX^2)/VIX^2$ for SPX and $ER = (\sqrt{RV} - VXO)/VXO$ for OEX. $LER = \ln(1+ER)$ denotes the continuously compounded return. Columns under "Mean, Std, Skewness, Kurtosis" denote the sample average, standard deviation, skewness, and excess kurtosis for the returns, respectively, based on daily data. "Auto" measures the monthly autocorrelation using non-overlapping data. We report the average of estimates at different starting dates. The last column under "Sharpe" measures the annualized Sharpe ratio of the excess return, defined as the mean monthly excess return over the Newey-West serial dependence adjusted standard deviation of the excess return, multiplied by the annualization scale of $\sqrt{12}$. We measure the Sharpe ratio on 30-day apart non-overlapping data and then take the average from the estimates with different starting dates.