The Nasdaq Volatility Index During and After the Bubble

David P. Simon

Dept of Finance
Bentley College
Waltham, MA 02452
dsimon@bentley.edu.
Tele: (781) 891-2489
Fax: (781) 891-2982

September 27, 2002

The author thanks the Hughey Center for Financial Services at Bentley College for the data and Bentley College for a summer research grant. The usual disclaimer applies.

The Nasdaq Volatility Index During and After the Bubble

Abstract

This paper examines the behavior of the Nasdaq 100 Volatility Index (VXN), constructed by the Chicago Board Options Exchange to reflect the implied volatility of near-term, at the money options on the Nasdaq 100 Index (NDX). The findings indicate that during both the technology stock bubble in the late 1990s and its aftermath, the VXN responds asymmetrically to positive and negative NDX returns, falling more in response to greater positive NDX returns and rising more in response to greater negative NDX returns. This finding is consistent with the VXN being a "fear index." The results also demonstrate that the VXN rises in response to technical indicators suggesting stronger trends, such as greater positive and negative deviations of the NDX from its 5-day moving average. In contrast to the behavior of the VXN, asymmetric GARCH models indicate that the conditional volatility of the underlying NDX returns does not fall in response to large positive NDX returns. The results also indicate that the VXN is a biased forecast of actual volatility, although it encompasses the information in asymmetric GARCH out of sample volatility forecasts. Nevertheless, the VXN averages 13 percentage points higher than actual volatility.

The Nasdaq Volatility Index During and After the Bubble

Studies by Fleming, Ostdiek and Whaley (1995) and Whaley (2000) examine the behavior of the Volatility Index (VIX), which is a measure of implied volatility calculated by the Chicago Board Options Exchange (CBOE) from near-term, close to the money options on the S&P 100 index. The VIX often is referred to as a "fear index" because its level indicates how much market participants are willing to pay in terms of implied volatility to hedge stock portfolios with S&P 100 index put options or to be long the stocks in the S&P 100 index with defined risk by buying call options. Evidence is consistent with this view, as Fleming et. al (1995) find that the VIX tends to rise following large selloffs and to fall following large rallies. Equity market strategists often interpret extreme levels of the VIX as contrary trading signals. High VIX levels indicate extreme pessimism, which causes equity prices to overshoot on the downside and leads to subsequent rallies.[1] Extreme low VIX levels reflect complacency among market participants, which sets up the market for disappointment and raises the odds of a market correction. Along these lines, McMillan (1996, p.32) states that "When options are too cheap, too many people are selling options at lower and lower prices. Contrary thinking tells us that such a market is about to explode."

Another major issue concerning the VIX, and implied volatility more generally, is whether implied volatility predicts actual volatility. While market participants commonly

[1] Simon and Wiggins (2001) find that when the VIX is in its top decile from January 1989 through June 1999, S&P futures are 2.8 percent higher on average over the next 30 days than they otherwise would be.

view implied volatility as a consensus forecast of actual volatility, evidence suggests that the predictive power of the implied volatility of S&P 100 index options is mixed. Canina and Figlewski (1993) examine the forecasting power of the implied volatility of S&P 100 index options for the actual volatility of S&P 100 index returns through option expiration from 1983 to 1987 and find that these forecasts are biased, and, more strikingly, typically do not have statistically significant forecasting power.[2] By contrast, Christensen and Prabhala (1998) examine the implied volatility of S&P 100 index options one month before expiration over a longer time span from 1983-1995 and find significant and unbiased predictive power**,** while Fleming (1998) finds significant, but biased, predictive power in the implied volatility of S&P 100 index options from 1985 to 1992. A major difference in these studies is that Canina and Figlewski (1993) examine daily overlapping forecasts, whereas Christensen and Prabhala (1998) and Fleming (1998) examine monthly non-overlapping implied volatility forecasts.[3]

Although the behavior of the VIX has been explored, no study to date has examined the behavior of the Nasdaq 100 Volatility Index, referred to as the VXN. The VXN is calculated by the CBOE in a manner similar to the VIX, based on the implied volatilities of at the money, near-term options on the Nasdaq 100 index (NDX). The present paper examines the empirical properties of the VXN from January 1995 through May 2002. The paper first models the time series properties of the VXN and focuses on

---

[2] In each of their subsamples, implied volatility fails the unbiasedness test, and in only 6 of 32 subsamples are the coefficients on implied volatility significantly greater than zero at the 5 percent level.
[3] Christensen and Prabhala (1998) show that in regressions of actual S&P 100 index volatility on implied volatility, the slope coefficients are considerably higher with non-overlapping data compared to daily overlapping data. They also point out that the extreme degree of autocorrelation in daily overlapping implied volatility forecast residuals raises the possibility of spurious regression problems, which causes regression coefficient estimates to be inconsistent. These issues are explored further in Hansen, Christensen and Prabhala (2001).

the extent to which implied volatility responds asymmetrically to NDX rallies and declines. Because the time series of the VXN includes both the technology stock bubble in the late 1990s and its aftermath, whether implied volatility responds differently to positive and negative NDX returns across these periods of extreme bullishness and bearishness is of interest. The paper also examines whether the VXN responds to technical indicators that suggest that the market is trending, such as deviations of the NDX from its 5-day moving average. The paper then examines the conditional volatility asymmetries of the underlying NDX index to determine whether they are consistent with the asymmetric reaction of the VXN to rallies and selloffs. Finally, the paper examines the predictive power of both the VXN and asymmetric GARCH out of sample forecasts for actual volatility and whether the VXN encompasses the information in GARCH forecasts.

The results indicate that the VXN responds highly asymmetrically to positive and negative NDX returns. Larger NDX selloffs lead to greater VXN increases, while larger NDX rallies lead to greater VXN declines. This asymmetry is extremely stable across the technology stock bubble and the post-bubble periods.[4] The results also demonstrate that the VXN is a biased forecast of actual volatility, although the VXN has significant predictive power for actual volatility and encompasses the information in out of sample time series forecasts. Nevertheless, the VXN overforecasts actual volatility on average by about 13 percentage points.

The paper proceeds as follows. The first section describes how the VXN is calculated, discusses the time series model of the VXN, provides preliminary statistical information concerning the variables in the model and presents estimation results of the model. The second section examines the time series properties of the volatility of the underlying NDX by estimating asymmetric GARCH models. The third section examines the forecast properties of the VXN and out of sample forecasts from asymmetric GARCH models. The final section summarizes the results of the paper.

I. The Nasdaq Volatility Index

I.A. Background Information on the VXN

The instrument underlying the Nasdaq Volatility Index (VXN) is the Nasdaq 100 Index (NDX). The NDX is a capitalization-weighted index comprised of 100 of the largest non-financial equities listed on the Nasdaq Stock Market. The NDX index was created in 1985 by the Chicago Board Options Exchange (CBOE) and is calculated and reported to market participants every 15 seconds during the trading day. Options traded on the NDX are European and may be exercised only on the last business day before the expiration date, which is the Saturday immediately following the third Friday of the expiration month. NDX options are cash settled based on the first five minutes of trading on the final trading day before expiration. NDX options trade for the three near-term months and for up to three additional months from the March quarterly schedule.

The CBOE began calculating and disseminating the VXN on an intra-day basis in January 1998 and has made the series available from the beginning of 1995. The closing VXN is based on quotes at the 4 p.m. EST close of regular hours trading on the Nasdaq. The VXN provides market participants with real-time information about the volatility assumptions embedded in the pricing of NDX options, based on the Black-Scholes model (1973) model with an adjustment for dividends. Similar to the more widely followed

[4] The terms "bubble period" and "post-bubble period" are used in this paper largely for expository purposes.

Volatility Index (VIX) based on S&P 100 index options, the VXN is constructed from the implied volatilities of eight options--the two closest to the money calls and puts for the two closest to expiration option contracts that have at least eight calendar days until expiration.

The VXN is calculated as follows: At the money implied volatilities are obtained for calls (puts) from each of the two closest to expiration months by interpolating between the implied volatilities of the two closest to the money calls (puts). These at the money implied volatilities for the two front call (put) contracts are interpolated to create implied volatilities for calls (puts) that have 22 business days before expiration. The VXN is formed by averaging these at the money implied volatilities constructed for calls and puts with 22 business days until expiration. The reason for using business days rather than calendar days for calculating the VXN is that, as shown by French and Roll (1986), stock return volatilities from Friday through Monday closes are similar to what they are from the close to close of consecutive trading days during the same week.[5] The VXN is calculated from the mid-points of the bid-ask spreads of the options rather than from the most recent transactions to avoid both bid-ask spread bounce and the possibility of stale option prices.

Figure 1 shows the daily history of the VXN and the NDX. The figure shows the massive and largely uninterrupted rally of the NDX from a level of 400 in January 1995 to a peak of 4700 in March 2000 and the selloff that subsequently took the NDX down to 1160 in May 2002. During the period from January 1995 through March 2000, the VXN moved mostly in a relatively narrow range between 20 percent and 40 percent, with one major departure to the upside in late 1998 associated with the Southeast Asian crisis, the

[5] If weekend days contribute to equity prices the same information flow and hence volatility as business days, the standard deviation of returns from Friday to Monday would be 3 times greater than the standard deviation of returns between consecutive trading days during the same week. The annualized standard deviations of daily returns over the

Russian default and the problems at Long Term Capital. Following the peak of the NDX in March 2000, the VXN moved in a range with a lower bound of about 40 percent and an upper bound of between 80 and 90 percent. This upper bound was reached during sharp NDX selloffs and following the World Trade Center Attack in September 2001.

I.B. A Time Series Model of the VXN

This section models the time series properties of the VXN. The model specifies the daily change of the VXN as a function of a constant, the lagged level of the VXN, separate variables for contemporaneous positive and negative NDX returns and technical indicators. The inclusion of the lagged level of the VXN allows for the possibility that the VXN reverts toward a typical level.[6] The specification also allows equal magnitude positive and negative NDX returns to have different effects on the VXN. Fleming et. al (1995) find that the VIX responds asymmetrically to positive and negative S&P 100 returns, with negative returns associated with implied volatility increases and positive returns associated with implied volatility decreases. This phenomenon has been attributed to the tendency of actual volatility to increase more following large negative returns than large positive returns of equal magnitude owing to financial leverage (see Black (1976) and Christie (1982)). However, French, Schwert and Stambaugh (1987) and Schwert (1989, 1990) demonstrate that this asymmetry is too strong to be explained by leverage alone.

---

sample period for consecutive days that are both trading days and for Fridays to Mondays are 27.7 and 28.3 percent, respectively.

[6] Whether mean reversion is different for positive and negative departures of the VXN from typical levels was also examined and asymmetries were not found.

The asymmetric reaction of the implied volatility of equity index options to positive and negative returns may stem from the ways that market participants use options to manage risk. One possible explanation is that market participants have a stronger inclination to buy put options to hedge portfolios following large market declines than following large rallies. Similarly, market declines may also give rise to a greater demand for the defined risk associated with buying call options rather than buying stocks, while rallies lead to reduced demand for the defined risk associated with buying call options rather than buying stocks. Either of these mechanisms--an increased demand for calls or puts--would have the same effect on implied volatility because put-call parity links the implied volatilities of calls and puts that have the same strike price and expiration date. For example, an investor who is long the NDX can convert this exposure to that of a long call position by purchasing put options or equivalently could sell the NDX and buy call options. Either would have the same effect on implied volatility. Although the asymmetric response to positive and negative returns has been documented in S&P 100 index options, such an effect has not been documented in Nasdaq index options.

This basic specification is augmented by variables used by technical analysts to gauge the strength of trends, which in turn may affect the demand for options and hence implied volatility.[7] If, for example, the NDX is expected to trend strongly upward, call option prices and their implied volatilities may be bid up by investors who prefer to buy calls rather than delta-equivalent stock positions. One explanation is that the presence of trends suggests a greater probability that the underlying NDX will be considerably higher or lower in the future, which would increase the potential payout of options and cause

[7] Technical indicators that suggest the emergence of trends could also induce trend followers to trade in the direction of the trend, which would be reinforcing.

implied volatility to be bid higher. Along, these lines, Leland (1996) demonstrates that investors who believe that market returns are less mean-reverting than the average investor will be inclined to buy options, whereas investors who believe that market returns are more mean-reverting than the average investor will be inclined to sell options. Also, because the delta of an option increases in absolute terms as its moneyness rises, a trending market would tend to cause options to outperform originally delta-equivalent stock positions, ceterus paribus.[8] One of the most commonly used indicators of trends is the deviation of the current level of the price of an instrument from a moving average of its recent levels.[9] In this paper, the percent deviation of the current level of the NDX index from its 5-day moving average is used, with separate variables for positive and negative deviations. If an increased demand for options causes implied volatility to rise as the strength of trends increases, the coefficient on positive and negative percent deviations from moving averages should be significantly positive and negative, respectively. By including both the contemporaneous return and the contemporaneous deviation of the NDX from its 5-day moving average, the model specification allows the impact of NDX returns on implied volatility to vary depending on the extent to which they lead to deviations of the NDX from its recent central tendency.

The model that is estimated is

[8] Delta is the first derivative of option prices in the Black-Scholes framework with respect to changes in the underlying instrument, and is the sensitivity of the price of an option to a small change in the underlying instrument. The tendency for deltas to rise in absolute value for calls (puts) as the underlying instrument rises (falls) is referred to as gamma, which is the second derivative of option prices to a change in the underlying instrument.

[9] Brock et al. (1992) find that moving average cross over trading rules are profitable in the Dow Jones Industrial average from 1897 to 1986.

$$\Delta VXN_t = \beta_0 + \beta_1 VNX_{t-1} + \beta_2\, NDXRET_t^{+} + \beta_3\, NDXRET_t^{-} \qquad (1)$$
$$+ \beta_4\, DEVMA5_t^{+} + \beta_5\, DEVMA5_t^{-} + å_t,$$

where the daily VXN change from the close at time t-1 to the close at time t is regressed on its lagged level, separate variables for positive and negative same day NDX returns and separate variables for positive and negative percentage deviations of the closing NDX from its 5-day moving average, respectively.

I. C. Background Data Information

Background information on the variables used in the model is shown in table 1. The table shows the means, standard deviations, top and bottom deciles and first four autocorrelations of the levels and daily first differences of the VXN, the daily NDX returns and the deviations of the NDX from its 5-day moving average. The VXN over the sample period averages 40.8 percent and has highest and lowest decile cutoffs of 63.8 and 25.6 percent, respectively. The daily first difference of the VXN averages .02 percentage points with highest and lowest decile cutoffs of 2.22 and –2.14 percent, respectively. Daily NDX returns over the sample period average .09 percent with highest and lowest decile cutoffs of 2.69 and -2.82 percent, respectively. Deviations of the NDX from its 5-day moving average have a mean of .12 percent with highest and lowest decile cutoffs of 2.9 and –2.9 percent, respectively. Autocorrelation is evident for the VXN and for the deviation of the

NDX from its 5-day moving average, but less so for the first difference of the VXN and the daily NDX returns.

I. D. The Model Estimation Results

The model estimates are shown in table 2. The model is estimated with the Newey-West correction for heteroscedasticity and autocorrelation at the first lag, as the autocorrelations of the residuals from the model show a spike at the first lag, but not at higher lags. The model is estimated over the whole sample period from January 3, 1995 through May 7, 2002 and for sub-periods corresponding to what is referred to as the bubble period from January 3, 1995 through the peak of the NDX on a closing basis on March 27, 2000 and what is referred to as the post-bubble period from March 28, 2000 through May 7, 2002.

The results for the whole sample period indicate that the VXN is mean reverting. A higher (lower) level of the VXN is associated with a subsequent decrease (increase) in the VXN that is statistically significant at the one percent level. The coefficient estimate of -.0227 on the lagged level of the VXN indicates that implied volatility shocks have half-lives of roughly 30 business days. The results also indicate that the impact of positive and negative NDX returns is highly asymmetric--larger positive returns are associated with greater VXN declines, while larger negative returns are associated with greater VXN increases.[10] Both coefficients are statistically significant at 1 percent confidence levels.

[10] The intercept coefficient of .006 and the -.57 coefficient on positive returns imply that positive NDX returns that are less than one percent are associated with small VXN increases, but NDX returns greater than 1 percent are associated with positive VXN changes. Thus, the statement that larger positive NDX returns are associated with

The coefficient estimates indicate that a 1 percent greater decrease in NDX returns is associated with a .55 percentage point greater VXN increase, whereas a 1 percent greater increase in NDX returns is associated with a .64 percentage point greater VXN decrease. T-tests not shown indicate that the absolute values of these coefficients are not significantly different from each other. These results imply that VXN changes respond equally but in opposite directions to positive and negative NDX returns.[11]

The results also indicate that both greater positive and negative deviations of the NDX from its 5-day moving average lead to statistically significant (at the one percent level) greater VXN increases. A one percent increase in both positive and negative deviations of the NDX from its 5-day moving average is associated with a roughly 1/5 percentage point greater increase in the VXN. Thus, despite the finding that positive NDX returns are associated with greater VXN declines, stronger positive NDX trends are associated with greater VXN increases. These findings can be reconciled along the following lines. Positive returns by themselves lead to a reduction in fear and a lower demand for protection from puts and a lower demand for the limited risk associated with being long calls rather than being long the NDX. However, to the extent that positive returns lead to positive deviations of the NDX from its 5-day moving average, the decrease in the NDX is mitigated. One possible explanation is that the trending behavior of the NDX leads to an increased demand for options because of gamma, which causes the deltas of the options to move in favor of call buyers.[12] In the case of negative returns that are associated

---

greater VXN declines is made for expositional ease and should be understood to include the possibility that greater VXN returns can also lead to smaller VXN increases.

[11] The asymmetry found here is greater than that found for the VIX. Fleming et al. (1995) find that the decline in the VIX for a positive S&P 100 index return is about half the increase in the VIX for an equal-magnitude negative S&P 100 index return using daily data from 1986-1992, excluding the stock market crash in October 1987.

[12] It is important to keep in mind the magnitudes of the effects. For example, if the NDX is equal to its 5-day moving average, a 1 percent NDX rally leads to a .55 percentage point drop in the VXN and the corresponding deviation from

with negative deviations of the NDX from its 5-day moving average, the downward trend of stock prices reinforces the effect of the negative return and the VXN rises by more.[13] Overall, the model explains about half of the variation of the daily VXN change.

The results for the sub-periods covering the bubble period from January 1995 through October 2000 and the post bubble period from October 2000 through May 2002 are similar to the overall results, to the extent that the VXN continues to respond highly asymmetrically to positive and negative NDX returns. Positive and negative returns in both sub-periods continue to have roughly equal but opposite effects on the VXN, and again the VXN is mean-reverting in both sub-periods. The effect of the technical indicators on the VXN across the two sub-periods varies. During the bubble period, greater positive but not greater negative deviations from the NDX 5-day moving average are associated with greater VXN increases. A possible explanation for greater negative deviations not having incremental effects on the VXN during the bubble period is that with the trend of the NDX overwhelmingly up, negative deviations from the 5-day moving average were viewed more skeptically as being the start of sustained downward trends. However, during the post-bubble period, both positive and negative deviations of the NDX from its 5-day MA lead to greater VXN increases.

Table 3 provides additional information about how the VXN responds to extreme NDX returns over the sample period from January 1995 to May 2002. Panel A shows that for the ten largest daily percentage NDX losses, which average –8.3 percent, the average VXN increase is 6.7 percentage points to an average level of 73.7 percent. As

---

the 5-day moving average leads to roughly a 1/5 percentage point increase in the VXN. In this case, the effect of the NDX rally on the VXN would be about half of what it otherwise would be.

shown in panel B, for the ten largest daily percentage NDX gains, which average 10.6 percent, the VXN falls an average of 6.4 percentage points to an average level of 70.5 percent. These VXN declines are fairly uniform--each of the ten largest percentage daily NDX rallies are associated with VXN decreases and the table shows that these declines typically occur from high VXN levels. These results are consistent with the VXN being a "fear index".

Overall, the results indicate that the VXN responds highly asymmetrically to large positive and negative NDX returns. A key question is whether this asymmetry reflects how the actual volatility of NDX returns responds to positive and negative returns or whether this asymmetry is driven by option trading dynamics, and, in particular, by fluctuations in the demand for the defined risk associated with buying options. The next section examines whether the conditional volatility of NDX returns is consistent with the asymmetry of the VXN with respect to positive and negative NDX returns.

II. Asymmetric GARCH Models of NDX Volatility

This section examines whether the actual volatility of NDX returns exhibits asymmetries with respect to positive and negative returns similar to those found in the VXN. The model estimated is the Glosten, Jaganathan and Runkle (GJR) model, which nests the GARCH model of Bollerslev (1986), and allows conditional volatility to respond

[13] While it is possible that the technical indicators proxy for very near-term volatility that is not somehow reflected in the model, including measures of recent volatility such as the most recent 5- or 10-day standard deviations of NDX returns does not affect the results.

differently to equal-magnitude positive and negative return innovations. The asymmetric GARCH model is specified as,

$$R_t = \mu + \eta_t, \quad \eta_t = N(0, h_t^2), \quad h_t^2 = \alpha_0 + \beta\, h^2_{t-1} + \delta_1 \eta^2_{t-1} + \delta_2\, D\eta^2_{t-1} \qquad (2)$$

where R_t is the daily return expressed as the difference from t-1 to t of the natural logs of the NDX, μ is the mean daily return, and η_t is the innovation of the mean daily return. The innovation is assumed to be normally distributed with a mean equal to zero and a conditional variance equal to h_t^2. The variable D is an indicator variable that takes on the value 1 if the lagged innovation is negative, and zero, otherwise. This specification allows equal-magnitude positive and negative squared lagged innovations to have different effects on the conditional variance, which is captured by the δ_2 coefficient. If this coefficient is significantly positive, negative return innovations have greater effects on conditional volatility than equal-magnitude positive returns. If the δ_1 coefficient is significantly positive, positive return innovations lead to conditional volatility increases, which would be inconsistent with the negative effect of positive NDX returns on implied volatility found in the previous section.[14]

Table 4 presents estimates of asymmetric GARCH models of daily NDX returns for the whole sample period from January 3, 1995 through May 7, 2002 and for the sub-samples corresponding to the bubble period from January 3, 1995 through March 27, 2000 and the post-bubble period from March 28, 2000 through May 7, 2002. For the whole period, the results indicate that both positive and negative return innovations give rise to significant conditional volatility increases. However, the coefficient estimates indicate that negative return innovations lead to roughly four times greater conditional volatility increases than positive return innovations of the same magnitude. The estimates also indicate that conditional volatility is mean reverting in both sub-periods. For the bubble period, the results indicate that both positive and negative return innovations give rise to higher conditional volatility, although the effect on conditional volatility from negative return innovations is about three times greater than from equal-magnitude positive return innovations. For the post-bubble period, the estimates indicate that positive return innovations no longer lead to statistically significant conditional volatility increases.

Overall, the results indicate that the conditional volatility of the underlying NDX index is far less asymmetric with respect to positive and negative returns than the VXN. In contrast to the evidence that the VXN tends to fall on large NDX rallies, no tendency is found for the conditional volatility of NDX returns to fall as positive returns rise. In the next section, the forecasting power of the VXN and asymmetric GARCH out of sample forecasts is examined.

III. The Predictive Power of VXN and GARCH Forecasts of Actual Volatility

III A. The Asymmetric GARCH Out of Sample Forecasts

[14] These models were also estimated with deviations of the NDX from its 5-day moving average in the conditional volatility equation, but these models did not converge.

This section examines the forecasting performance of the VXN and asymmetric GARCH out of sample forecasts (hereafter, GARCH forecasts) for the volatility of NDX returns from 22 business days before option expiration through option expiration. GARCH forecasts are obtained as follows: The asymmetric GARCH model described in the previous section is estimated with daily data over a pre-sample period from the beginning of January 1992 through the end of December 1994 to obtain initial parameter estimates. These estimated parameters are used to obtain one-day ahead volatility forecasts for 1995--the first year of the sample period--using the following formula,

$$E(h^2_{t+1}) = \alpha_0 + \beta h_t^2 + \delta_1 \eta_t^2 + \delta_2 D\eta_t^2. \qquad (3)$$

The conditional volatility forecasts more than one day ahead are formed using the following equation,

$$E(h^2_{t+s}) = \alpha_0 + (\beta + \delta_1 + .5^* \delta_2)\, E_t\,(h^2_{t+s-1}) \qquad \text{for } s > 1. \qquad (4)$$

Equations 3 and 4 reflect the assumption that positive and negative return innovations have asymmetric effects on conditional volatility only one day into the future. Because future positive and negative return innovations are assumed to be equally likely, more distant conditional volatility forecasts are invariant to the sign of the most recently observed innovation.[15] The forecast of the standard deviation of returns over the remaining life of

[15] If the asymmetry indicator variable is set equal to one because the most recent return innovation is negative, forecasting more than one step ahead using $\alpha_0 + (\beta+\delta_1+ \delta_2)\, E_t(h^2_{t+s-1})$ would be incorrect because doing so would assume a zero probability that future innovations will be positive. By the same reasoning, forecasting more than one step ahead following a positive innovation with the same formula absent the $\delta_2\alpha_1$ term would be incorrect because of

the nearby option contract from 22 business days before expiration is equal to the mean of the annualized daily standard deviation forecasts for each day through option expiration. For each subsequent year through the end of the sample period in May 2002, GARCH volatility forecasts are obtained in the same manner by re-estimating the models with the prior year of data added to the estimation period.

III B. The Data

Table 5 provides background information over the sample period from January 1995 through May 2002. The table shows the means, the standard errors and the bottom and top deciles of the actual volatilities, the forecasted volatilities and the forecast errors of the VXN and the asymmetric GARCH model. Random walk forecast errors also are included for perspective. The VXN is measured 22 business days before each of the monthly NDX option expirations, while actual volatilities and GARCH volatility forecasts are measured from 22 business days before option expiration through option expiration. The actual volatility is the sample standard deviation of NDX returns over the 22 business days before option expiration, with the mean return set equal to zero to avoid the potential bias that arises from using a noisy estimate of the true mean.[16]

The table indicates that the VXN averages 13-1/2 percentage points higher than actual volatility, while the GARCH volatility forecasts are on average only about 3 percentage points higher than actual volatility. The mean absolute forecast errors of the

the possibility that future innovations are negative. As a result, the asymmetry term is allowed to affect the forecast of the next day's volatility but it is assumed that positive and negative innovations have the same mean-reversion patterns and revert to the same unconditional mean. See Potter (1999) for a discussion of the complexities of forecasting multiple periods ahead with asymmetric GARCH models.

VXN are about twice as large as those of the out of sample GARCH volatility forecasts. The mean absolute error of the GARCH forecasts and the random walk forecasts are about equal, which indicates that the GARCH volatility forecast errors are roughly the same magnitude as the absolute changes of actual volatility. The VXN is more volatile than both actual volatility and the GARCH volatility forecasts, and has an upper decile cutoff of 64.88 percent, compared to 43.17 percent for actual volatility. The VXN also tends to be consistently higher than actual volatility, as the bottom decile of its forecast error is about 5 percentage points. In the next section, the predictive power and the bias of the VXN and the asymmetric GARCH out of sample forecasts are further examined.

IIIC. Predictive Power Tests

This section examines the predictive power of the VXN and the asymmetric GARCH out of sample volatility forecasts in both levels and in first differences over non-overlapping 22-business day horizons. In equation 5, the actual volatility of NDX returns is regressed on a constant and on the VXN and the out of sample asymmetric GARCH volatility forecast separately,

$$\sigma_{t,T} = \alpha + \beta \; \sigma_t^{VXN,G} + \varepsilon_{t,T} \tag{5}$$

$$\sigma_{t,T} - \sigma_{t-1,T-1} = \alpha + \beta \; [\sigma_t^{VXN,G} - \sigma_{t-1,T-1}] + \varepsilon_{t,T} \tag{6}$$

---

[16] See Figlewski (1997) for a discussion of this issue.

where $\sigma_{t,T}$ is the actual volatility of NDX returns from 22 business days before nearby option expiration through option expiration. σ_t^{VXN} and σ_t^{G} are the VXN and the GARCH forecasts of actual volatility, respectively, with both measured 22 business days before the nearby option contract expires. Equation 6 specifies the model in first differences, where the dependent variable is the change in actual volatility over the 22 business days before the previous option expiration to the 22 business days before the next option expiration, and $[\sigma_t^{VXN,G} - \sigma_{t-1}]$ are the predicted changes of actual volatility from its level before the previous expiration based on VXN and GARCH forecasts, respectively. If the volatility forecasts in equation 5 and 6 are unbiased, the intercept coefficients should not be significantly different from zero and the slope coefficients should not be significantly different from one. If the volatility forecasts have predictive power, the parameter estimates on the volatility forecasts should be significantly greater than zero.

Whether the VXN encompasses information in GJR volatility forecasts is examined by regressing actual volatility on a constant and on the VXN and on the GARCH volatility forecasts. These models also are estimated in both levels and in first differences,

$$\sigma_{t,T} = \alpha + \beta \sigma_t^{VXN} + \lambda \sigma_t^{G} + \varepsilon_{t,T}. \tag{7}$$

$$\sigma_{t,T} - \sigma_{t-1,T-1} = \alpha + \beta [\sigma_t^{VXN} - \sigma_{t-1,T-1}] + \lambda [\sigma_t^{G} - \sigma_{t-1,T-1}] + \varepsilon_{t,T}. \tag{8}$$

If the VXN encompasses the information in the GARCH volatility forecasts, the coefficients on the GARCH forecasts should not be statistically significant and the coefficient on the VXN should remain significantly positive. All of these regressions are estimated for the whole sample period with White's correction for heteroskedasticity.[17]

The estimates are shown in table 6. The results for the levels regressions indicate that both the VXN and the GARCH volatility forecasts have significant predictive power. However, both enter with coefficients that are significantly less than one, indicating that they are biased forecasts of actual volatility. In addition, the intercept coefficient is significantly positive in the GARCH regression, but is insignificant in the VXN model. These coefficient estimates indicate that the VXN tends to overforecast actual volatility over its entire range, whereas the GARCH forecasts tend to overpredict actual volatility over much of its range.[18] The VXN and the GARCH forecasts explain roughly 71 and 64 percent of the overall variation of the level of actual volatility, respectively. When both forecasts are included in the model, the VXN continues to enter with a significantly positive coefficient that is significantly less than one, whereas GARCH forecasts are no longer statistically significant.

When the models are estimated in first differences, both the VXN and the GARCH forecasts enter by themselves with significantly positive coefficients. However, these coefficients again are both significantly less than one, indicating that both forecasts are biased. The intercept term in the VXN model is significantly negative, whereas the intercept term in the GARCH model is not statistically significant. These coefficient estimates indicate that the VXN over its entire range tends to overpredict actual volatility changes, whereas GARCH forecasts tend to overpredict actual volatility changes in both

[17] These regressions are not run for sub-samples because with monthly data the post-bubble period has too few observations.

[18] The coefficient estimates indicate that the GARCH forecasts tend to underpredict actual volatility only when the GARCH forecast is less than .21, which, as shown in table 5, is not far from its lowest decile.

directions in that both positive and negative predicted volatility changes tend to overpredict actual volatility changes. These forecasts explain roughly 19 and 7 percent of the overall variation in actual volatility changes, respectively. When both the VXN and the GARCH forecasts are included in the first difference models, the VXN continues to enter with a significantly positive coefficient, and again the GARCH forecast does not enter the model with a significant coefficient. Overall, the results suggest that the VXN has significant predictive power in both levels and in first differences for actual volatility and encompasses the information in GARCH volatility forecasts, but is a biased forecast of actual volatility.

IV. Conclusion

This paper demonstrates that the Nasdaq Volatility Index (VXN) from January 1995 through May 2002 responds in an asymmetric manner to positive and negative NDX returns, rising more as NDX selloffs increase and falling more as NDX rallies increase. This asymmetry is stable across the bubble and post-bubble periods. However, this asymmetry is inconsistent with the conditional volatility of NDX returns, which shows no tendency to fall following positive return innovations. These findings suggest that the manner in which market participants use options to manage risk may explain some of the asymmetry. In particular, the asymmetry of the VXN is consistent with the demand for the protection associated with buying puts and the demand for the limited downside risk associated with buying calls increasing following NDX selloffs and falling following NDX rallies. The VXN also increases when the strength of trends increases, as measured by greater positive and negative deviations of the NDX index from its 5-day moving average.

This paper also demonstrates that the VXN has statistically significant predictive power for actual volatility, but is a biased forecast. In evaluating the predictive power of VXN forecasts, distinguishing between forecasting accuracy and predictive power is essential. Biased volatility forecasts may have substantial predictive, but may not provide reasonably accurate forecasts unless an adjustment is made for the bias. Although the VXN has significant predictive power for actual volatility and encompasses the information in GARCH volatility forecasts, its mean forecast error and mean absolute forecast error are 13-1/2 and 14 percentage points, respectively--roughly four times and two times greater than those of asymmetric GARCH out of sample volatility forecasts, respectively. These results indicate that Nasdaq volatility forecasts based on the VXN should be adjusted downward substantially.

Table 1. Summary Statistics of the VXN, NDX Returns and Technical Indicators
(Daily Data from January 3, 1995-May 7, 2002)

Variable	Mean	Std Dev	Top Decile	Bottom Decile	ρ_1	ρ_2	ρ_3	ρ_4
VXN	.4080	.1498	.6380	.2564	.99	.98	.97	.97
DVXN	.0002	.0206	.0222	-.0214	-.04	-.10	.01	-.06
NDXRET	.0009	.0240	.0269	-.0282	-.04	-.08	.01	.02
DEVMA5	.0012	.0248	.0285	-.0290	.63	.30	.08	-.04

VNX is the Nasdaq volatility index, calculated from at-the-money, near-term options on the Nasdaq 100 index, NDXRET are returns on the NDX index, and DEVMA5 is the percentage deviation of the NDX from its 5-day moving average.

Table 2. A Model of the Nasdaq 100 Volatility Index (VXN)

$$\Delta\, VNX_t = \beta_0 + \beta_1\, VNX_{t-1} + \beta_2\, NDXRET_t^{\,+} + \beta_3\, NDXRET_t^{\,-} + \beta_4\, DEVMA5_t^{\,+}$$
$$+ \beta_5\, DEVMA5_t^{\,-} + å_t,$$

The model regresses the daily change in the VXN on a constant, the lagged level of the VXN, separate variables for contemporaneous positive and negative returns on the underlying Nasdaq 100 index (NDX) and separate variables for positive and negative percentage deviations of the contemporaneous closing level of the NDX from its 5-day moving average. The model is estimated over the whole sample period and during the bubble period from the beginning of the sample from January 3, 1995 through March 27, 2000 when the NDX peaked and the post-bubble period from March 28, 2000 through the end of the sample period on May 8, 2002. The models are estimated with the Newey-West correction for autocorrelation at the first lag and heteroscedasticity. One and two asterisks denote statistical significance at the 5 and 1 percent levels, respectively.

Variables	Whole Sample Period (1/3/95–5/8/02)	Bubble Period (1/3/95-3/27/00)	Post-Bubble Period (3/28/00-5/8/02)
β_0	.0055**	.0024	.0122**
	(.0009)	(.0016)	(.0041)
β_1	-.0227**	-.0119*	-.0367**
	(.0036)	(.0053)	(.0091)
β_2	-.5668**	-.6831**	-.4724**
	(.0484)	(.0570)	(.0685)
β_3	-.6054**	-.8066**	-.4665**
	(.0553)	(.0658)	(.0882)
β_4	.2152**	.2664**	.1848*
	(.0489)	(.0430)	(.0882)
β_5	-.2201**	-.0176	-.3678**
	(.0539)	(.0752)	(.0758)
$\bar{R}^2$	.516	.535	.523

Table 3. Extreme Returns and VXN Changes

Panel A. Ten Biggest 1-Day NDX Percentage Losses

Date	$Return_t$	ΔVXN_t	VXN_t
08/31/98	-.0986	.0705	.5612
04/14/00	-.0980	.1423	.9317
01/02/01	-.0909	.0030	.8495
09/17/01	-.0825	.0942	.7326
12/20/00	-.0789	.1076	.8955
01/05/01	-.0781	.0940	.8276
04/03/01	-.0775	.0290	.7512
03/28/01	-.0769	.0177	.6837
10/27/97	-.0742	.1015	.4151
05/23/00	-.0739	.0081	.7202
Average	-.0830	.0668	.7368

Panel B. Ten Biggest 1-Day NDX Percentage Gains

Date	$Return_t$	ΔVXN_t	VXN_t
01/03/01	.1877	-.0896	.7599
12/05/00	.1168	-.0956	.6821
04/05/01	.1082	-.0008	.7629
04/17/00	.1010	-.1051	.8266
05/30/00	.1008	-.0422	.6776
04/18/01	.0953	-.0388	.6971
12/22/00	.0950	-.0634	.7923
10/13/00	.0910	-.0628	.5968
10/19/00	.0840	-.0827	.5426
04/25/00	.0799	-.0601	.7109
Average	.1060	-.0641	.70487

Table 4
An Asymmetric GARCH Model of the Daily Returns of the
Nasdaq 100 Index (NDX) from January 3, 1995 through May 7, 2002

$$R_t = \mu + \eta_t ,$$

$$\eta_t = N(0, h_t^2)$$

$$h_t^2 = \alpha_0 + \beta h^2_{t-1} + \delta_1 \eta^2_{t-1} + \delta_2 D \eta^2_{t-1}$$

	μ	α_0 x10^4	β	δ_1	δ_2	Likelihood Function
Jan. 3, 1995 -May 7, 2002	.0011* (.0004)	.0750** (.020)	.8939** (.0135)	.0359** (.0131)	.1145** (.0171)	6242
Jan. 3, 1995 -Mar. 27, 2000	.00167** (.0004)	.1831** (.0526)	.8202** (.0305)	.0636** (.0227)	.1267** (.0270)	4700
Mar. 28, 2000 -May 7, 2002	-.0033** (.0013)	.1218 (.1040)	.9287** (.0190)	-.0223 (.0142)	.1970** (.0378)	1553

R_t is the daily return on NDX index, μ is the average daily return, η_t is the innovation from the mean daily return, h_t^2 is the conditional variance of the innovation of daily returns, and D is an indicator variable that equals one when the lagged innovation of NDX returns is negative, and zero otherwise. Standard errors are in parentheses and one and two asterisks denote statistical significance at the five and one percent levels, respectively.

Table 5
Volatilities and Forecast Errors from 22 business days before expiration of the front NDX 100 option contract from January 3, 1995 – May 7, 2002 (NOBS=87)

Panel A: The Volatilities

	Mean	Standard Error	Bottom Decile	Top Decile
Actual Volatility	.2765	.1279	.1415	.4317
VXN	.4111	.1531	.1982	.6488
GARCH	.3058	.1485	.1786	.5491

Panel B: Forecast Errors

	Mean	Standard Error	Bottom Decile	Top Decile
VXN	.1346	.0821	.0535	.2331
GARCH	.0293	.0893	-.0674	.1279
Random Walk	.0050	.0904	-.0103	.1087

Panel C: Absolute Forecast Errors

	Mean	Standard Error	Bottom Decile	Top Decile
VXN	.1379	.0765	.0413	.2331
GARCH	.0668	.0658	.0057	.1474
Random Walk	.0651	.0626	.0082	.1352

Actual volatilities are the annualized standard deviations of the daily returns of the underlying NDX index from 22 days before nearby option expiration through expiration. The VXN is the Nasdaq 100 Volatility Index, which is calculated from the two closest to the money calls and puts for the two front contracts and interpolated to reflect the implied volatility of options that are at the money and have 22 business days until expiration. The GARCH volatility forecasts are out of sample volatility forecasts from 22 business days before option expiration through option expiration from out of sample asymmetric GARCH model estimates. The random walk absolute forecast errors are equal to the change in the actual volatility of the underlying NDX index return.

Table 6
Tests of the Predictive Power of the VXN and Out of Sample
GARCH Volatility Forecasts from January 1995-May 2002 (NOBS=87)

Levels: $\sigma_{t,T} = \alpha + \beta\ \sigma_t^{VXN,G} + \varepsilon_{t,T}$

$\sigma_{t,T} - \sigma_{t-1,T-1} = \alpha + \beta\ [\sigma_t^{VXN,G} - \sigma_{t-1,T-1}] + \varepsilon_{t,T}$

1st Differences: $\sigma_{t,T} = \alpha + \beta\, \sigma_t^{VXN} + \lambda \sigma_t^{G} + \varepsilon_{t,T.}$

$\sigma_{t,T} - \sigma_{t-1,T-1} = \alpha + \beta\, [\sigma_t^{VXN} - \sigma_{t-1,T-1}] + \lambda\, [\sigma_t^{G} - \sigma_{t-1}] + \varepsilon_{t,T}$

	Constant	VXN	GARCH	$\bar{R}$	D-W	Wald Tests
Levels						
	-.013 (.023)	.705** & (.064)	--	.71	2.06	385 (.00)
	.065** (.020)	--	.690** & (.075)	.64	2.13	20.12 (.00)
	-.009 (.023)	.644** & (.136)	.067 & (.149)	.71	2.11	--
First Differences						
	-.085** (.026)	.634** (.189)	--	.19	2.25	256 (.00)
	-.012 (.009)	--	.458* (.206)	.07	2.52	13.36 (.00)
	-.084** (.028)	.623** (.209)	.022 (.222)	.18	2.25	--

The returns volatility of the NDX index over the last 22 business days of the nearby option contract is regressed on a constant and on either the VXN, out of sample asymmetric GARCH volatility forecasts or both. The models are estimated in both levels and in first differences. Standard errors are estimated using the White correction for heteroskedasticity and are shown in parentheses. * and ** indicates significantly different from zero at the five and one percent level, respectively. & indicates significantly different from one at the five percent level. Wald tests are for the joint hypothesis that the intercept term is equal to zero and the coefficient on the volatility forecast is equal to one. P-values for the Wald tests are in parentheses.

References

Black, F. 1976, Studies of stock price volatility changes, Proceedings of the 1976 Meetings of the American Statistical Association, Business and Economics Section, 177-181.

Black, F. and M. Scholes, 1973, The Pricing of optionss and corporate liabilities, Journal of Political Economy 81, 637-659.

Bollerslev, T., 1986, Generalized autoregressive conditional heteroskedasticity, Journal of Econometrics 31, 307-327.

Canina, L.,and S. Figlewski, 1993, The information content of implied volatility, Review of Financial Studies 6, 659-681.

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

Christie, A., 1982, The stochastic behavior of common stock variances: value, leverage and interest rate effects, Journal of Financial Economics 10, 407-432.

Figlewski, S., 1997, Forecasting volatility, Financial Markets, Institutions and Instruments Vol.6, No.1. Boston MA: Blackwell Publishers.

Fleming, J., 1998, The quality of market volatility forecasts implied by S&P 100 index option prices, Journal of Empirical Finance 5, 317-345.

French, K. and R. Roll, 1986, Stock return variances: the arrival of information and the reaction of traders, Journal of Financial Economics 17, 5-26.

French, K., Schwert G. and R. Stambaugh, 1987, Expected stock returns and volatility, Journal of Financial Economics 19, 3-29.

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

Hansen, C., B. Christensen and N. Pbrabhala, 2001, Accounting for the overlapping data problem in the implied-realized volatility regression, Center for Analytical Finance Working Paper, University of Aarhus, Aarhus, Denmark.

Leland, H. 1996, Options and expectations, Journal of Portfolio Management, Special Issue, 43-51.

McMillan, L., 1996, The predictive power of options, Futures, 32-34.

Potter, S., 1999, Nonlinear time series modeling: an introduction, Journal of Economic Surveys 13, 505-528.

Schwert, G., 1989, Why does stock market volatility change over time?, Journal of Finance 44, 1115-1154.

Schwert. G., 1990, Stock volatility and the crash of '87, Review of Financial Studies 3, 77-102.

Simon, D. and R. Wiggins, 2001, S&P Futures Returns and Contrary Sentiment Indicators, Journal of Futures Markets 21, 447-462.

Whaley, R., 2000, The Investor Fear Gauge, Journal of Portfolio Management 26, 12-17.

Figure 1. The NDX and the VXN