The VOLATILITY-MARKET connection

Is everything you know about volatility wrong? Find out what history says about the volatility-market relationship — and what the VIX is saying about the stock market's 2004 prospects.

BY VICTOR NIEDERHOFFER AND LAUREL KENNER

Volatility is a crucial variable every market participant needs to consider. For speculators, volatility determines how much money to place on each trade relative to initial stake and stop point.

For investors, it determines how much to allocate between stocks and bonds, and how much to invest for a secure retirement.

For academics, volatility is one blade of the scissors in the fundamental theorem of finance — namely, that expected return is linearly related to volatility.

For the Spec Duo, volatility is a matter of life and death. On one day in 1997 we suffered the latter fate when volatility spiked by more than 50 percent, closing the market itself for the day.

In recent years, we have been one of the world's largest sellers of volatility. An article by Malcolm Gladwell in the April 22-29, 2002 issue of *The New Yorker* (see www.gladwell.com/2002/2002_04_29_a_blowingup.htm) featured a point-counterpoint between the derivatives expert Nassim Taleb, who "…buys out-of-the-money options by the truckload…[b]uys options on both sides, on the possibility of the market moving both up and down. And he doesn't bet on minor fluctuations in the market. Why bother?" and Vic, who likes to sell insurance to all those who believe financial Armageddon is around the corner and are willing to pay 100 times the fair price to receive it.

Taleb's idea is that no matter how many white swans have been seen, a black swan might still appear — and that such rare events are much more likely than people estimate. (A reference to 18th-century empirical philosopher and historian David Hume: "No amount of observations of white swans can allow the inference that all swans are white, but the observation of a single black swan is sufficient to refute that conclusion.")

Gladwell also wrote, "When [Taleb] contemplated the countless millions that Niederhoffer had made over the years, he could not escape the thought that it might all have been the result of sheer, dumb luck."

The rest of the article discusses how Taleb correctly predicted Niederhoffer would blow up because of his failure to take into account the inevitable volatility jumps Taleb foresaw. (Note: In 2003, the two funds Victor is currently associated with were up an average of 50 percent after fees, a result not inconsistent with the previous five years' performance.)

The shifting sands of volatility

In the good old days, volatility, which comes from a Latin word, *volare*, meaning "to fly," was traditionally used to describe emotions, especially those elicited by romance ("her spirit was volatile, but her heart tender…").

Today, the term most often refers to times of stock market panic and the irrational fear that prompts so many investors to sell or stay out of the market precisely when they should be buying. Typical in this regard was Zvi

continued on p. 30

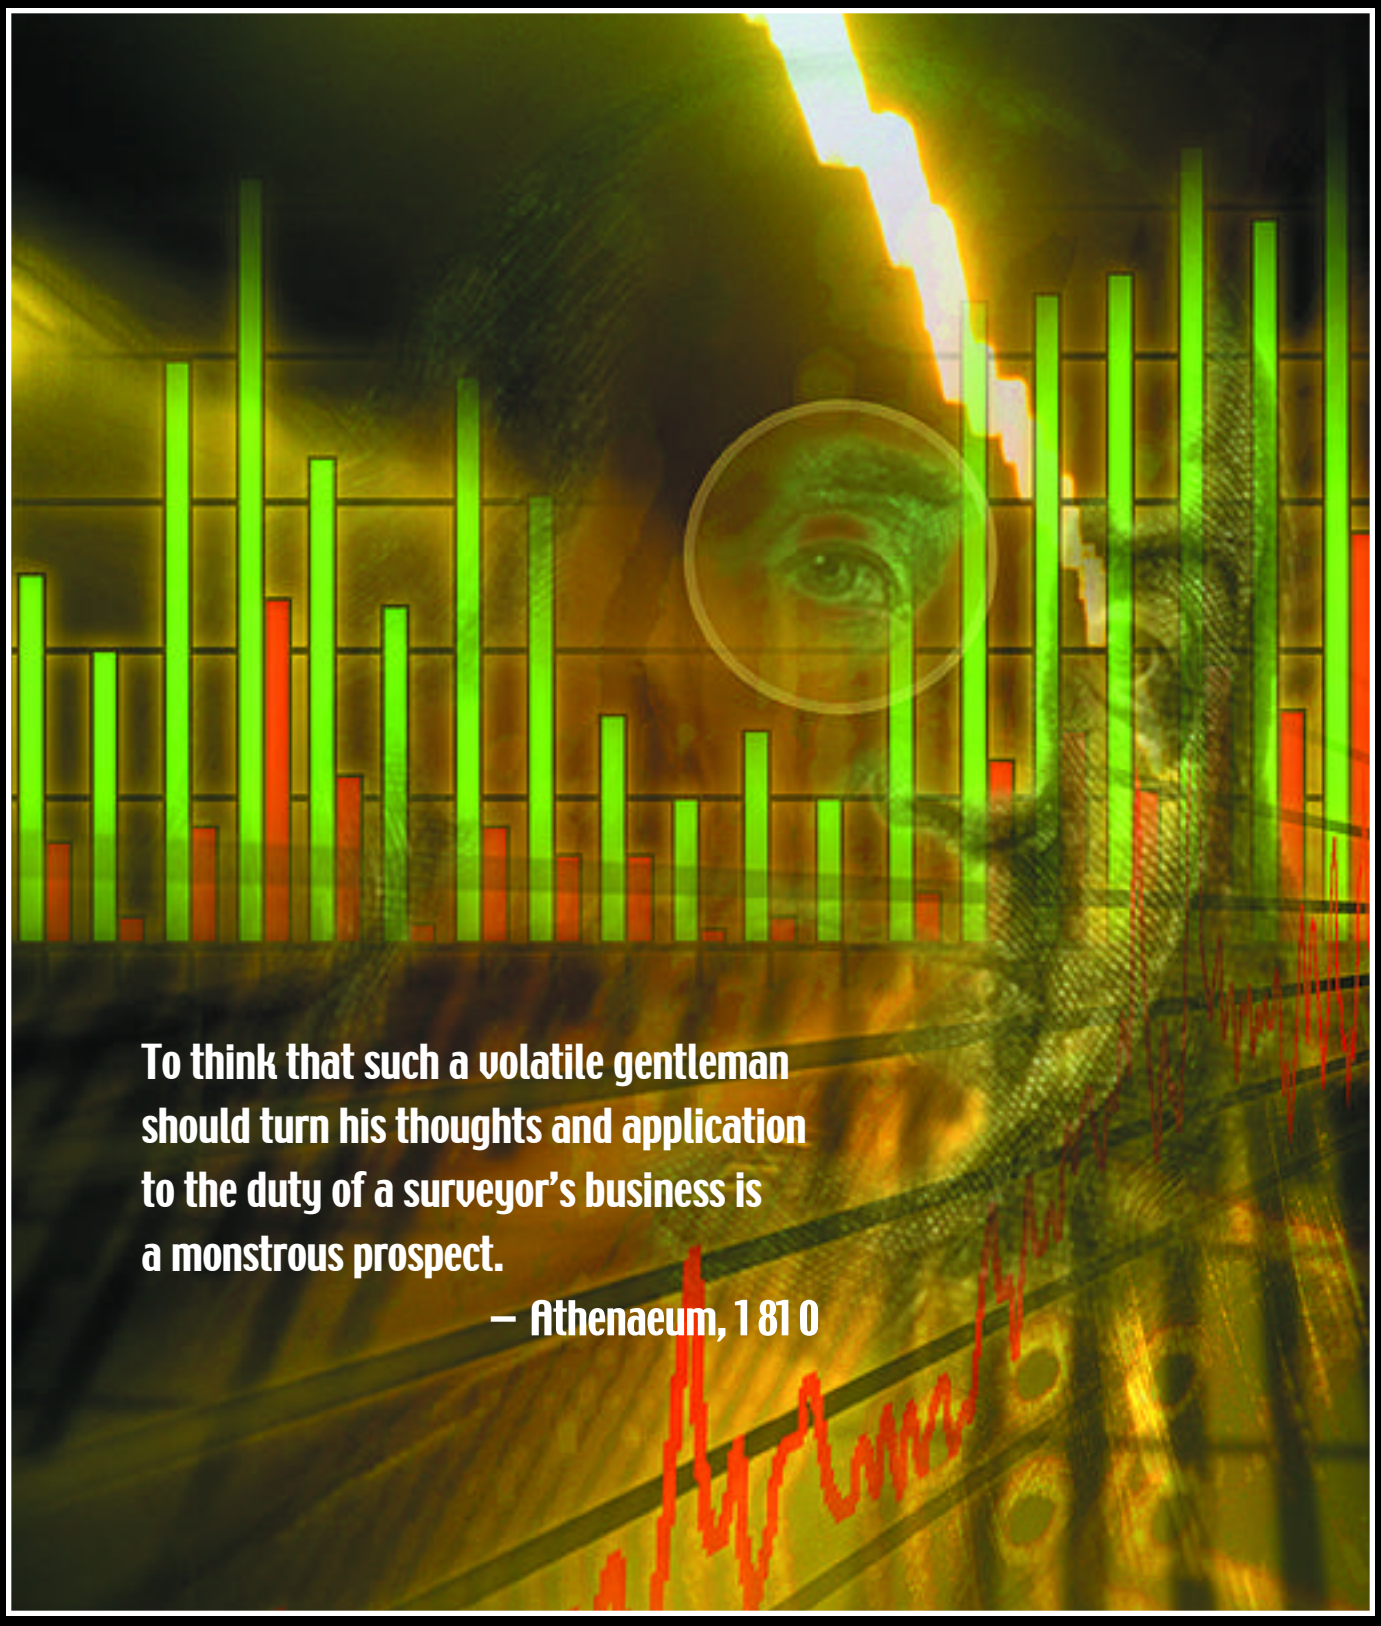
To think that such a volatile gentleman
should turn his thoughts and application
to the duty of a surveyor's business is
a monstrous prospect.
— Athenaeum, 1810

Bodie's widely acclaimed 2003 book, *Worry-Free Investing,* which advocated investors forego stocks altogether and instead buy inflation-indexed bonds.

Bodie's concern was the excessive volatility of stocks. And although it is true a widely accepted volatility measure, the Chicago Board Options Exchange's volatility index (VIX), registered 32 percent as of year-end 2002 — its highest year-end close since 1987 — investors who ignored Bodie's advice and bought stocks would have enjoyed average gains of more than 25 percent in 2003.

But now, an opposite phenomenon has emerged: By the end of 2003, stock market volatility had declined by almost half from year-end 2002 — its greatest yearly drop in 15 years. This decline came just as the Dow Jones Industrial Average broke above 10,000, the Nasdaq had climbed 51 percent for the year and all but one of the world's major stock markets were in the black.

Not since 1995 has there been such quietude. However, the calm of 1995 was followed by four consecutive years of gains, with the S&P returning more than 25 percent a year — despite the traditional association of low volatility with investor complacency and subsequent drastic declines.

Is the current volatility contraction the calm that precedes the storm? Or are we merely moving back to the good old days of the latter half of the '90s? Perhaps, we reasoned, in the process of studying the implications and causes of this volatility drop and sharing our conclusions with others, we might learn something ourselves that would be useful for the day and fray.

Tilting at windmills

An Internet search of "markets and volatility" returned 655,000 articles, including more than 1,000 scientific papers on the subject. However, with extensive academic and practical experience on the subject, we thought surveying this body of work would be somewhat Herculean but not impossible.

The fundamental theorem of finance is that returns are a linear function of risk — i.e., the higher the risk, the higher the expected return.

An extensive body of options pricing literature argues implied volatility is priced into the market and is, in general, much higher than actual, observed volatility. According to this research, options buyers pay a high price for "insurance," which represents a transfer of wealth from those who wish to reduce risk to those willing to assume it.

A conflicting school of thought has concluded implied volatility is not a good measure because it changes over time and is subject to unfathomable market dynamics —— the "black swan" problem (i.e., is it reasonable to bet that no black swans exist because you have never seen one?).

The more we read, the more we realized the task was not just Herculean, but Quixotic.

Rather than finding a general consensus about volatility we could apply to our business, we found an amorphous mass of conflicting conclusions, outdated data, over-determined results and theorizing without any relevance to market practicalities.

For example, it was claimed future returns and volatility were directly related to, and predicted by: the term structure of interest rates, inflation, money supply growth, dividends, cash flows, risk aversion, anchoring theory, the level of inflows to mutual funds, contagion, persistence, leverage, estimation error, the business cycle, stochasticity, regime shifts, consumption and GNP — indeed, all GNP *components*. (A good and relatively accessible review of these findings is contained in William Schwert's "Why Does Stock Market Volatility Change Over Time?" *Journal of Finance,* December 1989.)

Highlighting the confusing and conflicting conclusions about returns and volatility, Olesia Verchneko, in "The Determinants of Stock Market Volatility Dynamics," states:

> "A number of authors explored this [return-volatility] relation, but their conclusions are far from being uniform. Canina and Figlewski [1993], for example, find that implied and historical volatilities are virtually uncorrelated, and historical volatility, but not implied volatility, contains information about future return volatility. On the other hand, Christensen and Prabhala [1998] report that past volatility does not have much more explanatory power over implied volatility, and that implied volatility is an efficient (though biased) volatility forecast."

Similarly, St. Louis Federal Reserve economist Dr. Hu Guo, a leading expert in volatility research, in a review of the literature in his paper, "Understanding the Risk-Return Tradeoff in the Stock Market" (Federal Reserve Bank of St. Louis, January 2002), points out the cornerstone of modern finance theory is that risk-averse investors require compensation for any extra risk they bear in the stock market. However, Guo noted this "positive risk-return tradeoff has been argued to be inconsistent with data in several studies."

Guo went on to quote himself:

> "Interestingly, Guo shows the liquidity premium might be negatively related to the risk premium because conditional volatility is a u-shaped function of CAY [the consumption-wealth ratio]. Therefore we might fail to uncover a positive risk-return relation if we don't control for the effect of the liquidity premium on expected stock returns as observed in the post-World War II data."

Our reaction to the two quotes was similar to that experienced by Don Quixote on reading, in one of his favorite books on chivalry: "The reason of the unreasonableness which against my reason is wrought doth so weaken my reason that with all reason I do justly complain of your beauty."

According to *Don Quixote* author Miguel de Cervantes, it was this passage that pushed his protagonist over the edge and led to his noble attempt to become a latter-day knight-errant, "armed and mounted, in quest of adventures, redressing every species of error."

Our review of volatility literature broke our imagination in much the same way, and we embarked on a similar quest. We would develop our own meas-

continued on p. 32

FIGURE 1 VOLATILITY IN S&P 500 FUTURES, 1984-2003

The current volatility decline is reflected in the low number of one- and two-percent price moves in 2003.

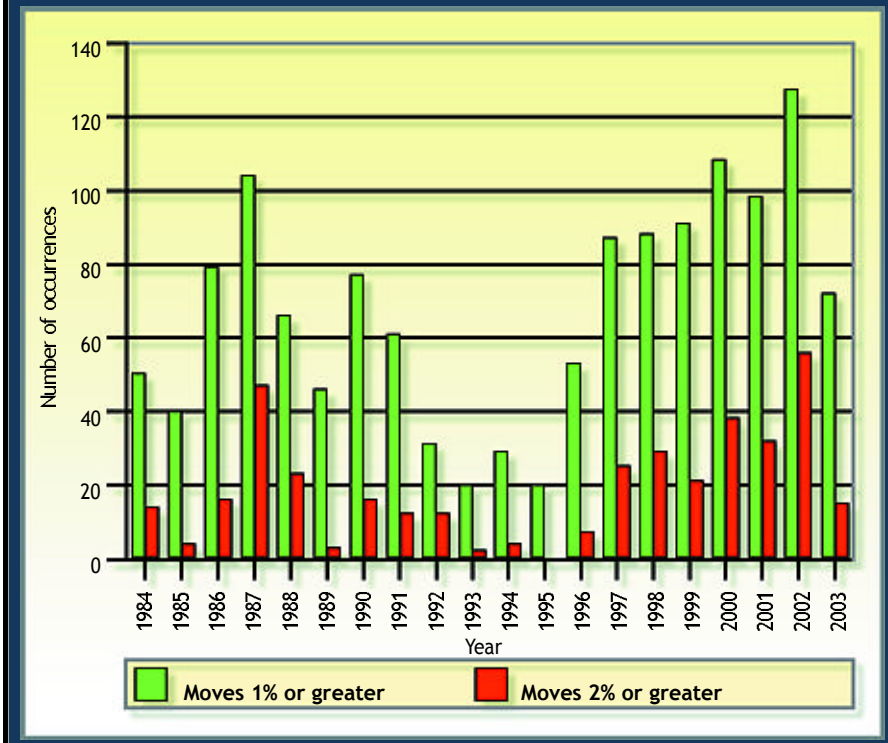

urements, collect our own data and perform our own calculations. We crunched the numbers four ways:

- We counted the daily changes of 1 percent and 2 percent in the S&P 500 index for each of the past 20 years.
- We measured the VIX over the past 15 years and quantified its relation to future yearly S&P changes.
- We developed a new way to measure the average risk of individual stocks and showed how this can be used to predict future stock returns.
- We calculated the actual volatility of the market between years and compared it to the forecasts of implied volatility for comparable dates.

The following sections describe our methods and conclusions in detail.

FIGURE 2 CBOE VOLATILITY INDEX (VIX)

The VIX, which measures the implied volatility of S&P 100 (OEX) options, is at its lowest levels since 1997.

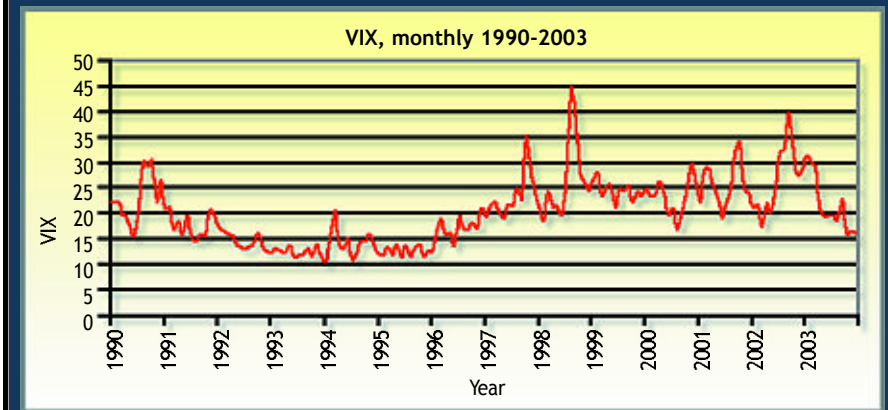

Source: Chicago Board Options Exchange

Volatility by hand

To get a handle on volatility, it's not necessary to know econometrics or dynamic programming. All it takes is a pencil and paper.

To begin, we calculated the number of 1-percent and 2-percent daily moves in S&P 500 futures in each of the last 20 years (see Figure 1, above left).

One striking feature of the chart is volatility in 2003 is the lowest since 1996. A mere 15 moves of 2 percent or more occurred in 2003, vs. an average of 32 per year over the previous six years.

Another clear indication is that a shift of some kind occurred in 1997. Before 1997, with the exceptions of 1987 and 1988, the norm was zero to 16 changes of 2 percent or more annually. But from 1997 on, each year has seen at least 20 such moves.

A reasonable question to ask is whether the low volatility of the pre-1997 years has returned, or whether 2003 was just an anomaly in the high-volatility, post-1996 world.

Interestingly, the numbers refute the common belief that large declines are more common than large rises. During the 20-year period, there were 183 rises of 2 percent or more and 177 declines of 2 percent of more — disproving the hoary adage that rallies from bottoms are sharper than falls from tops. (Doubtless dissenters from the grandstands will argue one-day moves are an inadequate basis of proof. To this, we suggest that statistics be put on the table that prove otherwise.)

The VIX rub

Market volatility is usually measured one of two ways. The first is the historical way illustrated in the previous section. But the market itself also contains within it a measure of volatility — it's the VIX, and it's designed to measure the type of volatility embodied in the S&P 100 (OEX) options that expire in the current and upcoming month.

The VIX has been computed since 1986, and is closely watched by many market participants. The typical interpretation is high VIX readings are bullish and low VIX readings are bearish. Several academic papers have attempted to prove this is the case and, indeed, in our worst-selling book *Practical Speculation,* we have a few pages devoted to VIX-based systems — buying, for example, when VIX breaks through 30 from below and closing out longs when it breaks through 25 from above.

A monthly chart of the VIX over the past 13 years shows that, at its late-2003 level of 15, the volatility index is the lowest it has been since 1996 (see Figure 2, opposite page). It also reflects the shifts noted in the 1-percent and 2-percent S&P moves in Figure 1. From 1997 forward, the median VIX value has been approximately 23 percent; before 1997, the median was about 16 percent.

A striking feature of both charts is the VIX tends to persist. The 13 to 17 range encompassed almost all the moves of the mid-1990s. And notice how in the new millennium levels in the 20s have occurred with inordinate frequency.

TABLE 1 S&P 500 INDEX RETURNS AND OEX VOLATILITY

Large declines in the VIX in one year tend to correlate to stock market losses the following year. Stock market variability is represented by the "AbsDev" column.

	1988-2003				
YrEnd	SP	YrChg	AbsDev	VIX	YrChg
2003	1,111.9	26.4%	15.0%	16.91	-45.3%
2002	879.8	-23.4%	34.7%	32.03	37.9%
2001	1,148.1	-13.0%	24.3%	23.22	-23.2%
2000	1,320.3	-10.1%	21.4%	30.23	13.2%
1999	1,469.3	19.5%	8.2%	26.71	5.1%
1998	1,229.2	26.7%	15.4%	25.41	2.1%
1997	970.4	31.0%	19.7%	24.89	14.9%
1996	740.7	20.3%	9.0%	21.67	56.0%
1995	615.9	34.1%	22.8%	13.89	3.3%
1994	459.3	-1.5%	12.8%	13.44	17.3%
1993	466.5	7.1%	4.2%	11.46	-15.4%
1992	435.7	4.5%	6.8%	13.55	-32.8%
1991	417.1	26.3%	15.0%	20.17	-14.4%
1990	330.2	-6.6%	17.9%	23.55	35.4%
1989	353.4	27.3%	16.0%	17.39	-6.2%
1988	277.7	12.4%	1.1%	18.53	-53.00%
Average		11.3%	15.3%	20.8	-0.3%
StdDev		17.9			

Waiting times

Another way to highlight the incredible volatility decline is to consider how much time has passed since the most recent big daily price drop. The waiting times between big declines makes interesting reading.

Only one decline of more than 25 points in S&P futures occurred in 2003, and that was on March 24 — exactly 119 trading days after the previous occasion, Sept. 27, 2002. That was a record waiting time for the previous six years, and 12 days longer than the previous record of 107 trading days before the 32-point decline of June 15, 1998. Since March 24, 2003, 194 trading days as of Dec. 31, 2003, have gone by without a 25-point decline — another record for "peace."

Big declines occurred sporadically within 10 to 20 days of each other throughout the 1997-2002 era. Then, in 2003, the declines suddenly dried up. It's the kind of thing statisticians study in survival data. After a certain period of exposure to disease, panic, product or idea, all the susceptible people in the population have already been infected, and the remainder have either died off or developed immunity.

(In passing, the analysis of waiting time, or lifetime, or survival statistics is a favorite pastime of ours, and one we have found highly useful in our market work. The best reference for studying it is *Analysis of Failure and Survival Data* by Peter J. Smith.)

Actual vs. implied volatility

One of the key questions of investing, both in theory and practice, is whether implied volatility is higher or lower than actual volatility. If it's higher, then options are priced too high. If it's lower, then options are priced too low.

The usual way to answer this question is to look at daily historical variability, adjust it upward to make it comparable to a longer-term volatility calculation, and then compare it to actual volatility.

A more direct way is to compute the actual longer-term variability of stocks on a historical basis and compare this with the average levels of longer-term implied volatility. Table 1 (above) does this, listing year-end VIX and S&Plevels from 1988 to 2003 along with measures of their average moves and variability.

Notice the S&P's average variability between years during this period was 18.2 percent. This compares to an average VIX level of 20.8 percent. However, the VIX shown is based on one- and two-month implied at-the-money options prices. These traditionally have been some 2 percent higher than year-end levels. So an adjustment of taking two points away from the VIX levels shown to adjust for year-end figures brings the average VIX level down to 19. Thus, implied volatilities and actual volatilities during this period have, on average, been roughly equal.

We can thus answer the central question of volatility with this simple table: at-the-money options are generally priced correctly. How predictable and deflating!

But that's guaranteed to happen, as

continued on p. 34

Vic's partner Steve Wisdom observed. Otherwise, one side would have all the other's money, and trade would cease.

Another interesting feature of the table is the price appreciation in the S&P in one year tends to move inversely relative to the previous year's VIX. In other words, low VIX levels at year-end tend to be associated with high price appreciation the next year. Thus, the four lowest years for the VIX were followed the next year by an average price appreciation of 15 percent. Conversely, the three years with the highest VIX readings were followed by an average move of 1 percent.

Overall, there was a -12-percent correlation between the VIX in one year and S&P appreciation in the next year. The facts do not support the theory that low VIX levels presage years of low price appreciation.

However, a very surprising result emerges when we analyzed *changes* in the VIX. These are quite predictive of S&P price appreciation the next year — much stronger than usually supposed. The correlation between the VIX change in one year and the S&P change the next year is a very strong 0.60. The more the VIX goes down, the greater the decline in the S&P next year. With 15 observations, a correlation this large would occur only three times in 100 by chance.

To make this a forecasting equation, we used regression analysis to come up with a predictive formula:

Next year's S&P return = 8 percent + (¼₀ * previous year's point change in VIX level)

For example, in the year 2002, the VIX level rose from 22 to 32, an increase of 10. The predicted 2003 S&P return would be 8 percent + (¼₀ * 10) = 33 percent.

Common sense and idiosyncratic risk

Given the low stock market volatility, what are the implications of this VIX-S&P forecasting model? The first implication follows from the basic theorem of modern finance: Expected excess returns of stocks are a linear function of non-diversifiable risk.

A variation of this idea is that future returns are linearly related to market variance. Both of these formulations can be found in most standard finance books.

These theories accord with common wisdom. For example, at the beginning of 2003, investors were highly fearful. They demanded higher returns to compensate them for higher risk. They got it in the form of such things as a 26-percent return in S&P 500 stocks, a 49-percent rise in the Nasdaq and similar gains in IPO stocks. Numerous systems buying low-priced stocks or beaten-down technology issues had returns in the triple digits.

Conversely, in 1999, investors were highly complacent and were willing to accept much hype from, for example, the technology-telecom-Internet sector. They saw little risk, and they were willing to accept a very low expected return — which is exactly what they got, with three consecutive years of declines totaling 38 percent in the S&P and 67 percent in the Nasdaq.

The only problem with these commonsense theories is they don't jibe with the facts. As mentioned above, the correlation between VIX and subsequent return is negative — i.e., low VIX values are followed by higher market returns.

Guo believes volatility does predict returns, but only if one takes into account the consumption/wealth ratio (a measure of liquidity) and the idiosyncratic risk associated with individual stocks.

We asked Guo for an explanation for why liquidity was important in predicting stock market returns (see "A word from Dr. Guo," right). His explanation is that investors demand a return for holding illiquid investments.

For example, when you hold a small stock you demand more of a return than if you hold a large stock because you can get out of the large one at any time without much of a hit from the last price.

We asked our friend Dr. Alex Castaldo, a trader and chief quantitative researcher at a large New York hedge fund, the same question.

"When the economy is in serious trouble — think of the 1930s — few of us can afford to invest in stocks," he says. "We need our labor income to cover our consumption and we are worried about unemployment. We have little liquidity. The few who invest experience enormous volatility, but they do so anyway in the expectation of big returns. The risk at this time is largely macroeconomic risk: Will the economy recover or not? Conversely, when the economy and the stock market look good, things calm down. More of us can invest and we are less worried about our jobs since our portfolios are larger relative to our consumption; some of us can even afford to retire early. Ironically, the expected return is less, precisely because plenty of people are willing to invest."

Regrettably, the consumption-wealth ratio is not reported contemporaneously, and the data is subject to many revisions. Therefore, another variable called "idiosyncratic risk" is used to measure liquidity. This variable is the part of a stock's variability that is not related to the market. Idiosyncratic risk is bad, Guo writes, because most people are not properly diversified.

Guo finds value-weighted idiosyncratic risk is a strong predictor of stock returns and, that when combined with VIX, it predicts very well how stocks are going to perform in the next year. High VIX levels and low idiosyncratic risk levels are good for subsequent returns, while low VIX levels and high idiosyncratic risk levels are unfavorable.

The only problem is it's hard to compute idiosyncratic risk. The standard method is to take the regressions of each stock on the NYSE, compute the residuals and then total them, weighted by market capitalization. That's the kind of information usually knowable a few years after the fact after some high-level computing with expensive data files.

We came up with a simpler method. The idiosyncratic risk in the market is directly related to how variable the stocks are relative to one another. A good measure of that variability is the range between the performance of, say, the best 50 and the worst 50 stocks in the S&P 500. For example, in 2002, the 50th-best stock was Sysco Corp., up 13.8 percent. The 50th worst stock was Comverse Technology, down 55.2 percent. The difference between these two figures — 69 percentage points, in this case — is a measure of how variable stocks are with respect to each other.

Table 2 (above) shows idiosyncratic risk (Range) and price appreciation for each of the last 15 years. Contrary to what you would expect, years when relative variability is greatest are followed by poor performance in the market, while quiet years are followed by outperformance. Indeed, the correlation between the idiosyncratic risk and the

TABLE 2 IDIOSYNCRATIC RISK AND S&P PERFORMANCE

The "Range" column shows the idio - syncratic risk, which is that part of a stocks variability that is not related to the market.

Year	Range	S&P return that year
2003	87.12	26.40
2002	68.89	-23.30
2001	77.52	-13.04
2000	118.10	-10.10
1999	160.60	19.50
1998	105.95	26.70
1997	76.20	31.00
1996	64.00	20.20
1995	70.00	34.10
1994	48.40	-1.50
1993	61.10	7.06
1992	54.90	4.40
1991	91.20	26.30
1990	46.20	-6.50
1989	58.90	27.20
1988	34.60	12.40
Avg.	76.40	11.30
St dev.	31.34	17.90

return the next year is a very large -0.48.

What does this indicate for 2004? As of year-end 2003, the 50th best performer in the S&P 500 was Xilinix, up 88 percent, while the 450th best was Fifth Third Bancorp, up 0.9 percent. The difference is 87 percentage points, 11 points wider than the average differential of 76 for the 16-year period. However, because the ranges we reported during this period were not adjusted for survivor bias (the shift in favor of those stocks that remained at the end of an observation period vs. those that disappeared), we may assume the range should be adjusted down by five or 10 percentage points. Thus, idiosyncratic variability is probably some 18 percent greater than normal in 2003. By itself, this is a relatively small differential without much predictive significance.

However, the studies did uncover one strikingly negative feature of year-end 2003 variability. Recall the VIX change in one year is highly correlated (at the 60-percent level) with S&P returns the following year. The VIX had fallen in 2003 by 16 percentage points, its largest decline since 1988 — and this is very negative for 2004 stock market returns.

Using the regression formula given in the "Actual vs. Implied Volatility" section, the predicted S&P 500 change for 2004 would be: 8 percent – (¼ * 16) = -32 percent.

This highly bearish prediction is based on a regression that has an r-squared (a measure of how well a regression line fits the actual data points) of 35 percent. While the regression relation itself is very unlikely to have arisen by chance, individual predictions based upon it have an extremely wide margin of error and are highly uncertain. And while the negative prediction is somewhat deflating to the bullish case for years after major declines in VIX, we place no undue reliance on the forecast

The Don Quixote syndrome

At the end of his quest to redress the wrongs of the world, save damsels in distress and relive the heroic deeds glorified in the great books on chivalry, Don Quixote is filled with sadness. He believed the books had led him to folly and disaster, so he decided to burn them and disown his heir if she married anyone who had ever even heard of chivalry.

We felt a similar melancholy at the conclusion of our own quest to unlock the implications of volatility. The usual ideas contained in the textbooks are either wrong or misleading. The central niche we fill in the market — i.e., the selling of options — is neither good nor bad, but rather mediocre. Worst of all, we are heavy-hearted about the likely rocky journey in store for stocks in 2004.

Volatility is a concept that resonates in both markets and life. As Dr. Brett Steenbarger, an expert in behavioral finance and author of *The Psychology of Trading: Tools and Techniques for Minding the Markets*, noted, "Like a quarrelsome couple that finds its best lovemaking after a blowout argument, investors get their sweetest love from the Market Mistress after her most tempestuous outbursts."

For information on the author see p. 10.

Thanks to Duncan Coker and Steve Wisdom for their help in preparing this article.

A word from Dr. Guo

Dr. Hu Guo is one of the most respected and prolific authors in volatility research. Reading some of his articles sparked our quest. We thought it would be fair to Dr. Guo and informative for readers to get his feedback on the results of our quest. He responded as follows:

"I received a Ph.D. in economics from NYU and my dissertation is about the equity-premium puzzle in a dynamic model. I found that, in addition to the risk premium as in standard CAPM, investors demand a 'liquidity premium' as well, because of limited stock market participation.

"If we use CAY (the consumption/wealth ratio) as a proxy for the liquidity premium, as my model suggests, I find that realized volatility and CAY are strong predictors of stock market returns. I am quite excited about this result because the early authors actually find a weak or even negative risk-return relation. I also address the risk-return trade-off in a paper [authored] jointly with R. Whitelaw at Stern School.

"Many people talk about the idiosyncratic volatility (IV) because most investors are poorly diversified. Goyal and Santa Clara (2003, Journal of Finance) is the first paper that address the predictability of IV for stock market returns using time series data. We start investigating whether IV provides additional information beyond CAY and market volatility and we find more interesting results that value-weighted IV is a strong predictor of stock returns.

"In a few other papers, I, or with my co-author, address the implication of stock-return predictability on the cross section of stock returns.

"I would say that my research provides some new insights on the way we think about the equity market."

Guo's papers are available at http://research.stlouisfed.org/econ/guo/.