THE JOURNAL OF DERIVATIVES

Editor's Letter

Stephen Figlewski

JOD 2017, 25 (1) 1-3
doi: https://doi.org/10.3905/jod.2017.25.1.001
http://jod.iijournals.com/content/25/1/1

Institutional Investor Journals
1120 Avenue of the Americas, 6th floor,
New York, NY 10036, Phone: +1 212-224-3589

THE JOURNAL OF DERIVATIVES

VOLUME 25, NUMBER 1 FALL 2017

Did volatility die, or is it just on vacation? As I am writing this on August 8, 2017, the VIX index stands at 9.67. Since July 1, it has closed below 10 on 11 out of 25 trading days. To see how unusual this is, in the prior 27 years in which the VIX has been calculated (6,675 trading days), it only closed below 10 on nine occasions. During that period, the VIX averaged 19.8. I have written previously about how confusing this is, given the visibly high level of overall uncertainty in the world. Obviously, the stock market is unworried, though, because it has hit new all-time highs over and over. Many commentators are now warning that stock prices will very soon be in bubble territory, if they aren't there already. Perhaps we should also begin to worry about a matching anti-bubble in volatility.

Concern that the VIX may be too low goes beyond the general thought that stock market volatility will probably revert toward more normal levels eventually. In the last few years, a profusion of financial products that are tied to volatility, and to the VIX in particular, have been launched. Exchange-traded funds and notes allow investors to take a position on volatility, both direct and inverse, with leveraged exposure up to three times that of the index itself. Because the VIX index is not investible (it is more like a temperature than a stock index in that regard), these structured products use futures contracts to gain exposure to the VIX. As volatility has drifted downward during 2017, products with inverse exposure have profited handsomely. How much lower can the VIX fall before it rebounds? How sharp might the rebound be if the anti-bubble bursts? How much exposure on the short side is in the form of written options, such that trades to rebalance delta hedges will require buying volatility when it is spiking upward? Most importantly, how badly might investors with levered short positions be hurt in the process? One big difference between the VIX and a regular stock index like the S&P 500 is that the VIX can move extremely rapidly without much trading, simply because expectations or risk tolerance change. One cannot imagine the S&P index doubling in a matter of days, but it is quite possible for the VIX to go from, say, 10 to 20 almost overnight.

Based on the publicity and discussion one sees on the web and elsewhere, it is not unreasonable to feel that many investors in VIX-based exchange-traded funds and notes do not fully understand how they behave and, in particular, how their returns are related to the underlying stock market. First, the stated objective of the majority of such funds is to match the *one-day* return on VIX futures, levered up by the fund's leverage factor. This introduces a convexity effect that hurts returns for positions held over longer periods and means that incorporating volatility products into a buy-and-hold stock portfolio as a way to diversify and to manage volatility risk over time is a mistake. These products should only be used for very short-term strategies (as their sponsors emphasize).

Second, and I believe less well understood, is that there is a big difference between the behavior of the VIX index and a 30-day futures contract on the VIX. Because the VIX itself cannot be bought and held, the futures price is not determined by a cost-of-carry relationship. Instead, futures prices should be based on where the market expects the VIX index to be on the maturity dates of those contracts. The connection between what the spot VIX does and changes in VIX futures prices does not arise from arbitrage, so it is weak (the correlation is well below 1.0 even for the nearest future). Moreover, if the futures market is informationally efficient and expectations are unbiased, VIX futures prices should not incorporate any expected drift over time, either up or down. If the two-month future is higher than the one-month future, the difference reflects the market's rational forecast that implied volatilities will be that much higher in two months than in one month; the spread is no more likely to narrow over time than to widen. One might, and probably should, wish to modify the pure expectations pricing model by adding suitable risk premiums, but how this works out in practice is not simple. Rather than offering a full treatment here, I will simply suggest that many investors may be taking positions in these products without a full appreciation of either the risk exposure or the expected returns.

P.S. (24 hours later): Oops! After Trump threatened North Korea with "fire and fury," the S&P 500 opened down about 0.5%, but the VIX jumped up to 12.31, an increase of about 27%.

While awaiting future developments in this arena, let us turn our attention to the current issue of *The Journal of Derivatives*.

A typical derivatives pricing model begins with a set of assumptions about the returns distribution of the underlying asset and the market environment and derives a single arbitrage-free price. In the real world, however, transaction costs and parameter uncertainty drive a wedge between prices that are really too low and those that are really too high. In the range between these bounds, the price is OK in the sense that investors don't think it is so wrong that they are confident a trade against the market will be profitable. Mathematically, such a range forms a scalable convex set called a *cone*. In our lead article, Madan and Schoutens develop a theory of conic option pricing, in which the pricing cone is determined by deforming the expected returns distribution both upward and downward, and the option value is assumed to be the price at which buying and selling quantities balance out. Their general approach can be applied with a broad set of probability models. In the next article, Lu and Xu find a wrinkle in a newish technology for option valuation in lattice models. The standard numerical approximation techniques, trees and Monte Carlo simulation, try to model the instantaneous behavior of the underlying factor(s) that cumulates to the final payoff on the derivative at expiration. Improving accuracy requires shortening the time step, which increases the total number of calculations exponentially. In contrast, the willow tree method uses many prices but only a small number of time steps, with many long branches that connect each price node to the full set of prices at the next time step. Earlier research has shown that this structure, when applicable, can produce very accurate answers in much less execution time. This article's main contribution is showing how the willow tree can be efficiently extended to handle a multivariate problem: pricing a convertible bond that is exposed to both stock price and interest rate risk.

In 1977, Miller first raised the idea that constraints on short selling can introduce bias into stock prices. Investors with favorable opinions about a stock can buy it, but those who believe it is overpriced rarely sell it short, so their unfavorable information does not get into the market price on an equal basis with the bullish opinion. Research over the years has largely supported this theory, but the effect has been mitigated by the introduction of puts and calls, which provide alternative ways to take a short position in a stock. Yet the connection among differences of opinion, short sales, and future realized returns is still a little unclear. Short sales are not treated symmetrically with purchases, but they aren't banned either. If a stock has an unusually high level of short interest, does this indicate substantial adverse opinion that investors are expressing by taking short positions, or does it mean the impediments to shorting are simply smaller for that stock than for others (in which case, high short interest would have no informational content)? Du finds that the case is really the former and that the overvaluation effect is strongest when both market beliefs are heterogeneous and the cost of shorting is high. Next, Uhrig-Homburg and Unger present a new approach to valuing a gas storage facility, given the strong seasonal pattern in the price of natural gas, the nature of consumption as a flow, gas futures and forward contracts that apply not to a single date but to the cost of gas supplied over a period of time, and finally the facility operator's ability to increase returns by exploiting short-run volatility. The trick is to fit a market model to the data but to construct a continuous forward curve from it.

In the following article, Chin and Garriott report on a Penny Pilot experiment on options at the Montreal Exchange, in which the minimum tick size was reduced from $0.05 to $0.01 for options on five groups of stocks at different times. Their results indicate that bid–ask spreads narrowed and liquidity increased, especially for low-priced out-of-the-money contracts. Finally, Kan takes a fresh look at the problem posed by basket options, in which the payoff is determined by a portfolio of lognormally distributed securities that produce a nonlognormal returns process for the portfolio. He applies a statistical trick that greatly accelerates the calculation, such that the technique is basically as accurate as the most accurate alternative model and as fast as the fastest one.

It is now time for your editor to enjoy what remains of the summer. By the time you are reading this, the fall semester will be upon us. May I offer my best wishes that it will prove successful for us all.

Stephen Figlewski
Editor