The Carlyle Group

Structural Relationships and Optimal Portfolio Construction

By Jason M. Thomas

September 2014

Executive Summary

- Traditional portfolio optimization depends on investors' ability to accurately forecast future returns, variances, and correlations across assets.

- This is a virtually impossible task because returns and correlations change abruptly in ways that cannot be known in advance.

- Instead of relying on inherently frail *statistical* relationships, investors should base allocation decisions on the enduring *structural* relationship between risk and return.

- Direct investments in implied volatility indexes like the VIX and related contracts can provide a perfect "hedge" because increases in risk "feedback" into the price of the asset; when the conditional volatility of an asset increases, its price falls.

- Empirical results align with theory, as increases in the front-month VIX futures contract have translated to downward adjustment in stock prices. Data demonstrate that the VIX is a reliable proxy for the market-wide price of risk and is generally able to explain stock returns, changes in credit spreads, and anticipate "flight to safety" episodes six months in advance.

- Alternative hedge assets, like government bonds, are generally poor substitutes for direct investments in volatility because their returns are positively correlated with stocks in certain circumstances. It would not be surprising if current conditions caused the positive relationship between stocks and bonds to emerge on a sustained basis over the next several years.

1 The Failure of Existing Portfolio Allocation Methods

Efficient portfolio construction depends on diversification, which requires offsetting deviations in returns. If some asset prices rise in circumstances when others' prices decline, then combining those two sets of assets into a single portfolio will result in less volatility and higher returns per unit of risk than those assets could generate in isolation. The simplicity of the objective masks the difficulty of finding assets whose returns *reliably* deviate from those of the market as a whole in all future states of the world.

If expected returns and correlations were static, portfolio construction would be a straightforward exercise: investors could compile past data on the mean, variance, and correlation of various securities' returns to calculate the "optimal" portfolio based on the investor's risk tolerance. Unfortunately, the evidence is overwhelming that this is not the case. The "price of risk" varies through time in ways that result in profound shifts in expected returns and correlations across asset classes.[1] Weights that would have been optimal over the past five or ten years could generate inordinate losses over the next five.

When correlative relationships shift, the offsetting deviations counted on to reduce portfolio volatility often fail to materialize. This phenomenon is especially apparent during periods of market stress, when the correlations of disparate asset classes rise in ways that could not have been predicted based on past observations. As shown in Figure 1, the average trailing twelve month cross-asset correlation between emerging market currencies, high-grade corporate credit, commodity prices, and the S&P 500 rose from -2% in the period preceding the crisis to a high of over 70% in the period following the Lehman bankruptcy filing. A spike in equity market risk infected most other investible asset categories, causing many "well diversified" portfolios to suffer losses as though filled entirely with stocks.

Many investors were not protected from sizeable losses in the crisis because of the failure of the existing portfolio construction paradigm. Traditional portfolio construction can only succeed to the extent that assets' returns, correlations, and volatilities can be known in advance. Instead of questioning whether this is a realistic objective, researchers and practitioners responded to past failures by building more and more sophisticated statistical models. The main breakthrough had been the introduction of conditional volatility models that allow the volatility of returns to vary through time.[2] While this class of models can describe past data quite well, they generally rest on the assumption that relationships will eventually return to "normal" and therefore fail when the shift in returns or correlations is of a more permanent character ("structural breaks").[3]

Figure 1: Cross-Asset Correlation (Currencies, Stocks, Commodities, Credit)[4]

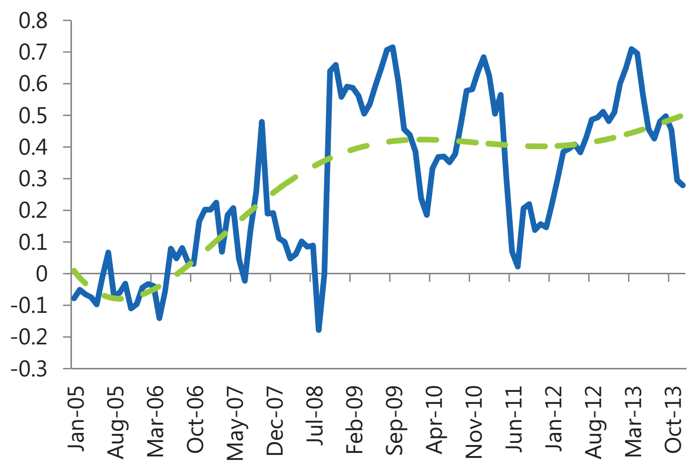

For example, conditional volatility models were first adapted to a portfolio setting by treating individual assets' returns as a linear function of a small(er) number of risk factors.[5] These "factor models" generally work well until the old factors estimated from past data lose their influence and other factors emerge unexpectedly.[6] Researchers and practitioners have addressed structural breaks by taking averages of parameter estimates through time, changing the frequency at which returns data are collected, and introducing new techniques more robust to out-of-sample data.[7] Unfortunately, none of these models have represented the hoped-for breakthrough for portfolio strategy.[8]

The past approach has failed to yield a durable solution because it is largely based on a false premise. The problem is not the technology used to make the forecasts; it's the conceit that inherently random returns processes can be reliably forecast. What's required is not a better model, but a fundamental rethink of the portfolio construction problem. Instead of relying on inherently frail statistical relationships, investors should base allocation decisions

1 The "price of risk" is defined as the expected return per unit of risk as measured by an asset's sensitivity to the stochastic discount factor, or pricing kernel, which measures representative agent's marginal rate of substitution of wealth across future states of the world. C.f. Cochrane, J. (2012), "Discount Rates," *Journal of Finance*.

2 Bollerslev, T., Engle, R., and Wooldridge, (1988), "A Capital Asset Pricing Model with Time-Varying Covariances," *Journal of Political Economy.*

3 Rapach, D. and Strauss, J. (2008), "Structural breaks and GARCH models of exchange rate volatility," *Journal of Applied Econometrics.*

4 Data from Federal Reserve Bank of St. Louis

5 Engle, R. et al. (1990), "Asset pricing with a factor ARCH covariance structure: Empirical estimates for treasury bills," *Journal of Econometrics.*

6 Engle, R., Sheppard, K. (2001), "Theoretical and empirical properties of dynamic conditional correlation multivariate GARCH," NBER working paper 8554.

7 Jagannathan, R. and Ma, T. (2003), "Risk reduction in large portfolios: Why imposing the wrong constrain helps." *Journal of Finance.*

8 Gourierouxa C., et al. (2009), "The Wishart Autoregressive process of multivariate stochastic volatility," Journal of Econometrics.

on the enduring structural relationship between risk and return.

2 The Structural Relationship Between Risk and Return

Asset pricing theories differ in many important respects, but virtually all share one common theme: an asset's expected return increases with its risk.[9] In most cases, the expected return is captured by the mean (expectation) or first moment of the distribution of future returns and the risk is the standard deviation (volatility) or second moment of the same distribution.[10] This structural relationship predicts that if risk (volatility) increases, so too must expected returns.

The positive relationship between volatility and future (expected) returns generates a negative relationship between volatility and contemporaneous returns. That's because increases in risk "feedback" into the price of the asset; when the conditional volatility of an asset increases, its price falls.[11] Increases in conditional volatility reflect an increase in the perceived riskiness of the asset or portfolio, which causes its price to drop and generates the observed negative returns. (The same works in reverse, as declines in risk perceptions cause the price of the asset to increase).[12]

For example, assume the current market price of a $100 payment due in one year is $90 (11% expected return). If the risk of nonpayment increases, the expected (future) return on the asset must increase. With the promised payment fixed, the only way for the expected return to increase would be for its $90 present value to decline today. If we assume the price fell to $80, the increase in volatility would generate an 11% decline in returns today but a 14% increase in expected returns (a 25% return from $80 to $100).

The major driver of return processes is not changes in fundamentals, but changes in the compensation investors require to hold different assets. In the past, changes in asset prices were thought to reflect market expectations for future fundamentals: an increasing stock market signaled faster growth; a sell-off meant "the market" expected growth to slow or contract.[13] This theory was disproven by empirical research demonstrating that asset price volatility is too great to be explained by changes in expectations of future cash flows. It is now widely accepted that asset prices adjust in response to perceptions of risk and these adjustments account for the bulk of the observed volatility in stock returns.[14] As in our example, price declines generally do not arise because of a fall in expected future cash flows, but because investors demand more compensation for bearing the risk that those cash flows may not materialize.

Forward-looking volatility, therefore, represents a perfect "hedge" asset because it captures changes in risk perceptions or the "price of risk."[15] As a result of the structural relationship between risk and return, implied volatility is associated with declines in asset prices in the current period (negative contemporaneous correlation) as well as increases in future (expected) returns. Therefore direct investments in conditional volatility should reliably offset declines elsewhere in the portfolio while simultaneously providing more resources to invest in risky assets when expected returns are highest.

3 VIX as a Market Measure of Conditional Volatility

In 1993, the Chicago Board Options Exchange (CBOE) introduced the CBOE Volatility Index (VIX Index).[16] The first VIX used option prices to measure the 30-day implied volatility of the S&P 100 index. In 2004, volatility became an investible asset class when the CBOE updated the calculation of the VIX index and introduced trading of VIX futures contracts. VIX options began trading two years later.

The VIX (post-2004) captures the 30-day standard deviation of the distribution of future returns of the S&P 500 implied by the prices of a portfolio of options written on the index (SPX). Conceptually, implied volatility is similar to a bond's yield to maturity.[17] A yield to maturity is the discount rate implied by the relationship between the current price of a bond and its promised future cash flows; the VIX is the level of volatility implied by the relationship between options at various strike prices that mature in the future and the current price of the S&P 500.[18]

The VIX has been called the "fear index." That is because the options used to calculate the VIX are mainly used by hedgers who buy index puts to protect against a downward

9 Engle, R. et al. (1990).

10 This is generalizable to virtually all asset pricing theories if we exclude labor income and treat the distribution of returns on aggregate wealth as the source of all future consumption opportunities.

11 In the 1970s, academics believed the observed relationship between volatility and returns was due to the "leverage effect," where the volatility of firm assets is constant, with equity volatility changing in response to shifts in effective leverage due to variation in stock returns. This theory has generally been dismissed as academics have come to appreciate the time series variation in discount rates described in Cochrane (2012). Empirical evidence has confirmed that most of the volatility-returns relationship is due to time-varying risk premia, including studies that focus on the value of firm assets and hence abstract from capital structure. C.f. Bekaert, G. and Wu, G. (2000), "Asymmetric Volatility and Risk in Equity Markets," *Review of Financial Studies*; Choi, J. and Richardson, M. (2012), "The Volatility of the Firm's Assets," American Finance Association 2010 Meetings; and Duffee, G. (1995), "Stock Returns and Volatility, A Firm-Level Analysis," *Journal of Financial Economics*.

12 Note that in real time the price decline of the stock can occur in advance of any change in market-based measure of volatility.

13 Cochrane (2012).

14 C.f. Campbell, J. (1991), "A variance decomposition for stock returns. *Economic Journal*.

15 The reference to hedge asset comes from Merton, R. (1973), "An Intertemporal Capital Asset Pricing Model," *Econometrica*.

16 Prior to

17 Whaley, R. (2008), "Understanding VIX," Working Paper, Owen Graduate School of Management, Vanderbilt University. Whaley is credited as the founder of the VIX

18 Technically, the VIX is the standard deviation of future returns that would be expected to obtain if investors were risk neutral. As a result, the VIX reflects the expected volatility under the physical probability measure plus a time-varying risk compensation component.

move in stock prices.[19] The VIX is therefore analogous to the price of insurance, which adjusts to account for changes in perceptions of the probability of a "bad" event or the capacity of other market participants to provide protection against it. The VIX and its futures contracts are considered the cleanest representation of the price of risk available to market participants.

4 Empirical Performance of Front-Month VIX Futures Contracts

Since its introduction in 2004, the performance of the front-month VIX futures contract has been as strong as one would anticipate from a structural rather than statistical relationship. Increases in the front-month futures contract reliably feedback into stocks, lowering their price and generating the anticipated negative correlation between implied volatility and contemporaneous returns. At the same time, the value of the front-month futures contract and its subsequent changes explain trailing twelve month (TTM) returns on the S&P 500 (i.e. variation in returns reflects observed variation in the VIX contracts). Increases in the front-month contract also have reliably forecast changes in credit spreads as well as declines in short-term interest rates arising from a "flight to safety." Both findings emphasize the manner in which implied volatility can be generalized across asset classes to improve portfolio construction.

Figure 2 plots the front-month VIX futures contract against the inverse of the level of the S&P 500. The close correlation between the series (73.6%) over nearly 11 years of observations indicates that stock prices have reliably adjusted downward in response to an increase in implied volatility. Likewise, declines in the front-month VIX futures contract have been associated with sustained increase in stock prices.

Figure 2: VIX Futures and the Inverse of S&P 500[20]

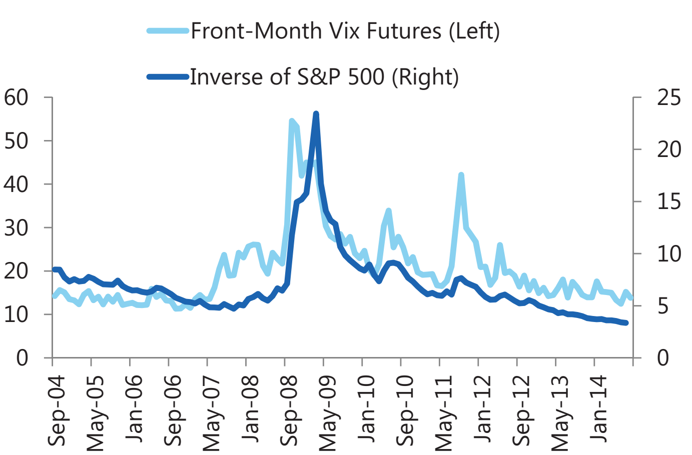

Since implied volatility is thought to be a measure of risk, Figure 3 looks at the relationship between the front-month VIX futures price and stock market valuations. The futures contract is converted to a discount factor applied to expected earnings on the broad stock market. If option-implied volatility is a reliable proxy for the price or risk, increases in the contract should increase expected returns by reducing the price an investors is willing to pay for a given stream of earnings. This is exactly what is observed in Figure 3: since 2004, the VIX discount factor has been 77.5% correlated with variation in price-to-earnings ratios. As in the $100 payment example, increases in the VIX reduce market prices relative to earnings (lower P/E ratios) so as to increase expected returns. Recent research documents that these market-wide relationships are also observed at the individual stock level, when implied volatility is measured relative to returns implied by analyst price targets.[21] As shown in Figure 4, the observed variation in the VIX front-month futures contract explains 63% of TTM returns on the S&P 500.

Figure 3: VIX Futures and S&P 1500 PE Ratios[22]

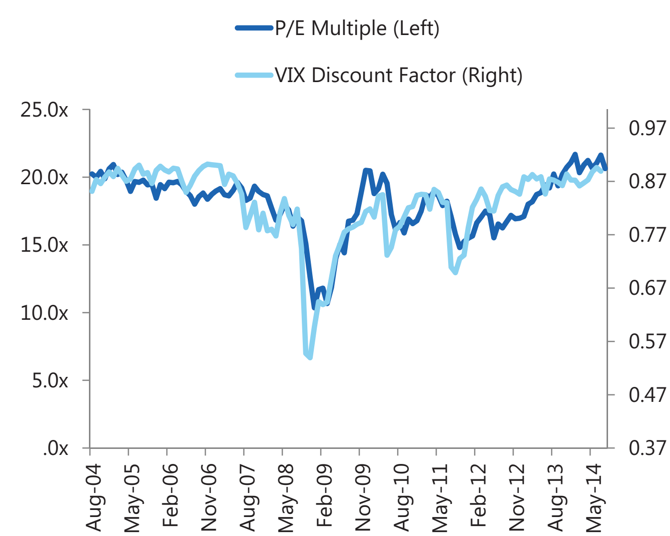

The use of the VIX of a tool to measure expected returns extends beyond equity markets. Structural models of credit risk treat creditors (bondholders) as short a put option on the assets of the firm, with a strike price equal to the face value of debt.[23] In this framework, the credit spread is equal to the put option's premium, which is positively related to the implied volatility of the assets. According to this model, the implied volatility of a broad market index, like the S&P 500, should be closely related to the aggregate price of credit risk.

19 Whaley (2008).

20 Data from Federal Reserve Bank of St. Louis and Bloomberg (UX1 Index).

21 Bali, et al. (2014), "Option Implied Volatility, Skewness, and Kurtosis in the Cross-Section of Expected Stock Returns," Working Paper, McDonough School of Business.

22 Data from Federal Reserve Bank of St. Louis and Bloomberg (UX1 Index).

23 Merton, R.C. (1974), "The Risk Structure of Interest Rates," *Journal of Finance*.

Figure 4: VIX Futures and TTM S&P 500 Returns[24]

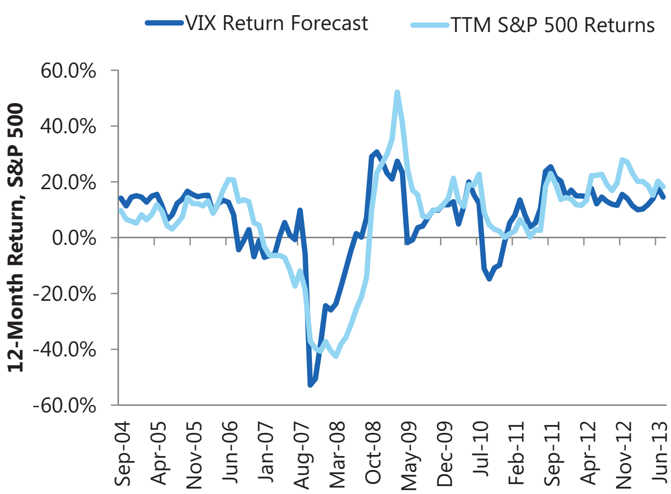

As shown in Figure 5, the data bear this out: since 2004, changes in the front-month VIX futures contract have explained 84.7% of the monthly variation in the option-adjusted spread on B-rated corporate credit. That the correlation is so high despite the low overlap between the two indexes is strong evidence that the VIX corresponds to a market-wide price of risk rather than security-specific factors. Direct investments in front-month VIX futures provided a nearly perfect hedge for the increase in credit spreads during the initial money market disruption of 2007, the spike in spreads in the fall of 2008, and the increase following the U.S. rating downgrade of 2011.

Figure 5: VIX Futures and Speculative Grade Credit Spreads[25]

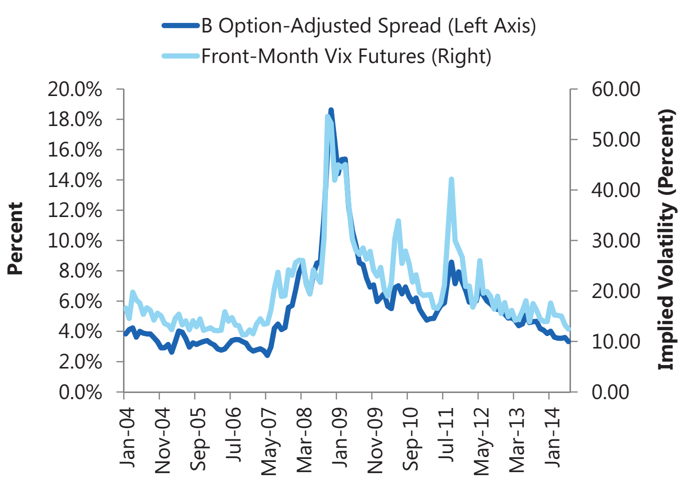

Finally, since implied volatility proxies for the price of risk, increases in the VIX have generally preceded "flight to quality" episodes. As shown in Figure 6, increases in the front-month VIX futures contract from 13 to 17 have been associated with a decline in 90-day Treasury yields from 3.1% to 1.24%, on average, over the next six months. The increase in the price of risk causes investors to shift portfolios into less risky assets, leading to cumulative declines in short-term U.S. interest rates.

Figure 6: VIX and 90-Day Treasury Yield in 6 Months[26]

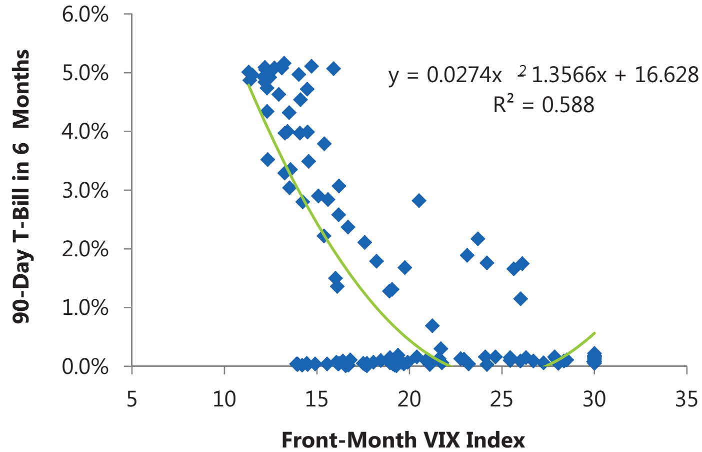

5 Relying on Bonds as a Hedge Asset is Dangerous

It is common to describe the VIX as "negatively correlated" with stock returns. While technically true, the observed negative relationship is just the empirical manifestation of a deeper, structural relationship between risk and return. Yet, the language of correlation confounds this deeper association and could lead many investors to mistakenly conclude that direct investments in volatility is just one of many hedge assets "negatively correlated" with market returns.

For example, many portfolios purporting to offer high returns per unit of risk today rely on leveraged positions in government bonds to hedge equity market risk. Longer-dated bonds have indeed been negatively correlated with stocks in recent years and provided a spectacular hedge in the second half of 2008 when an unleveraged position in 10-year Treasury notes gained 17% as the S&P 500 declined by 33%.[27] Yet, Treasury bonds have performed unevenly as a hedge asset through the course of history. In the 1970s and 1980s, Treasury bonds added to investors' systematic risk by moving in the same direction as stocks and the broader economy.[28] It would not be surprising if current conditions caused this positive relationship to again emerge over the next several years.

24 Data from Federal Reserve Bank of St. Louis and Bloomberg (UX1 Index). Note that this is not a trading strategy as the evolution of the VIX cannot be observed in advance. It simply demonstrates that the current level of the VIX and its evolution explain most of the observed returns in S&P 500 over the same period.

25 Data from Federal Reserve Bank of St. Louis.

26 Data from Federal Reserve Bank of St. Louis. Since nominal rates cannot decline below zero, VIX index levels consistent with negative interest rates have been censored.

27 Thomas, J. (2013), "Escaping the Fed's Asset Price Cycle," The Carlyle Group, July 2013.

28 Campbell, J.Y. et al. (2014), "Monetary Policy Drivers of Bond and Equity Risks," NBER Working Paper No. 20070.

ting losses elsewhere in the portfolio, bonds were actually intensifying them.

Figure 7: Market Beta of a 10-Year Treasury Portfolio[29]

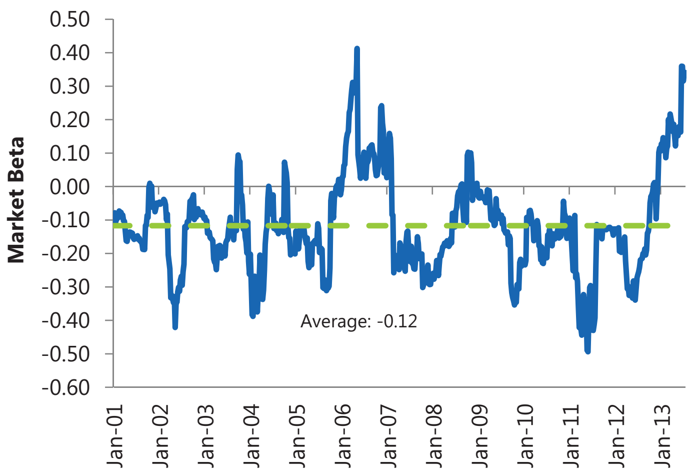

An asset's risk can be decomposed into two categories: fundamental, which involves uncertainty associated with future sales, earnings, and cash flows; and financial, which concerns the appropriate rate at which those future cash flows are discounted to present value.[30] For volatility-based hedge investments, the origin of the risk is of no consequence: the increases in risk and risk perceptions are reflected in implied volatility measures and feedback into asset prices. Conversely, bonds provide a hedge only for fundamental risk. Increases in financial uncertainty also exert downward pressure on Treasury bonds by raising the "term premium" investors demand for bearing the risk that interest rates will rise in the future.

The "taper tantrum" was the clearest manifestation of the limited value bonds can play as a hedge asset when financial risk increases. On May 22, 2014, Fed Chairman Bernanke suggested that if the labor market continued to improve, the Fed could begin to "taper" its asset purchases in the next few meetings. This information increased financial risk by increasing the likelihood that Fed policy rates would rise sooner than expected, which would increase the discount rates applied to all future cash flows. Over the next month, the prices of virtually all financial assets adjusted downward, with long-duration Treasury notes actually dropping by even more than the S&P 500 (-8% relative to -5%).[31] As shown in Figure 7, during the course of 2013, the market beta of the 10 year Treasury moved from -0.3 to 0.3. This means that instead of offsetting losses elsewhere in the portfolio, bonds were actually intensifying them.

Given the low levels of interest rates and uncertainty about the pace at which the Fed and other central banks will remove accommodation, it seems reasonable to think the "taper tantrum" will not be an isolated incident. Valuation ratios are high today largely because of low expected discount rates. Should the Fed have to raise rates faster than currently contemplated, bond prices would likely fall in tandem with stocks to reflect the new path for short rates and also to provide compensation for the risk that rates will increase more in the future.

6 Conclusion

Traditional portfolio construction strategies depend on advance knowledge of returns, variances, and correlations over the relevant investment window. Since there is no way of securing such information, it is no surprise that these techniques have largely failed to protect investors from sizeable losses. Instead of relying on inherently unstable statistical relationships, investors should base portfolio allocation decisions on the enduring structural relationship between risk and return. Implied volatility indexes like the VIX and related contracts can provide perfect hedge assets. Implied volatility is associated with declines in asset prices in the current period (negative contemporaneous correlation) as well as increases in future (expected) returns.

As a proxy for the market price of risk, VIX can also be used to hedge credit risk exposure and forecast "flights to safety." Other candidates to serve as a portfolio hedge for macro risk, like government bonds, are poor substitutes for direct investment in volatility. Bonds only move in the opposite direction as stocks in certain contexts, as seen by the positive correlation between bond and stock returns in the 1970s and 1980s. If current conditions caused this positive relationship to again emerge over the next several years, the advantages associated with direct investments in volatility could be even greater.

29 Bank of America Merrill Lynch and Federal Reserve Board of Governors. The beta is measured as the weekly return on the 10-year Treasury portfolio net of the 90-day bill rate regressed on the excess weekly return of the value-weighted U.S. stock market.

30 Campbell, J. and Shiller, R. (1988), "The Dividend Price Ratio and Expectations of Future Dividends and Discount Factors," *Review of Financial Studies.*

31 S&P Capital IQ Database Accessed July 23, 2013. The "Long Duration Treasury" is a portfolio with a weighted average effective duration of 15 years.

Jason Thomas is a Managing Director and Director of Research at The Carlyle Group, focusing on economic and statistical analysis of the Carlyle portfolio, asset prices, and broader trends in the global economy. Mr. Thomas is based in Washington, D.C.

Mr. Thomas' research helps to identify new investment opportunities, advance strategic initiatives and corporate development, and support Carlyle investors.

Mr. Thomas received a B.A. from Claremont McKenna College and an M.S. and Ph.D. in finance from George Washington University where he was a Bank of America Foundation, Leo and Lillian Goodwin, and School of Business Fellow.

Mr. Thomas has earned the Chartered Financial Analyst (CFA) designation and is a financial risk manager (FRM) certified by the Global Association of Risk Professionals.

Contact Information

Jason Thomas

Director of Research

jason.thomas@carlyle.com

(202) 729-5420

THE CARLYLE GROUP

GLOBAL ALTERNATIVE ASSET MANAGEMENT