Predicting Volatility

Stephen Marra, CFA, Senior Vice President, Portfolio Manager/Analyst

Uncertainty is inherent in every financial model. It is driven by changing fundamentals, human psychology, and the manner in which the markets discount potential future states of the macroeconomic environment. While defining uncertainty in financial markets can quickly escalate into philosophical discussions, volatility is widely accepted as a practical measure of risk. Most market variables remain largely unpredictable, but volatility has certain characteristics that can increase the accuracy of its forecasted values. The statistical nature of volatility is one of the main catalysts behind the emergence of volatility targeting and risk parity strategies.

Volatility forecasting has important implications for all investors focused on risk-adjusted returns, especially those that employ asset allocation, risk parity, and volatility targeting strategies. An understanding of the different approaches used to forecast volatility and the implications of their assumptions and dependencies provides a robust framework for the process of risk budgeting.

In this paper, we will examine the art and science of volatility prediction, the characteristics which make it a fruitful endeavor, and the effectiveness as well as the pros and cons of different methods of predicting volatility.

Introduction: Statistical Properties of Volatility

"The starting point for every financial model is the uncertainty facing investors, and the substance of every financial model involves the impact of uncertainty on the behavior of investors and, ultimately, on market prices. The very existence of financial economics as a discipline is predicated on uncertainty."[1]

A big part of risk management, asset allocation, and trading in financial markets is quantifying the potential loss of assets. In order to measure these potential losses and make sound investment decisions, investors must estimate risks.

Volatility is the purest measure of risk in financial markets and consequently has become the expected price of uncertainty. The trade-off between return and risk is critical for all investment decisions. Inaccurate volatility estimates can leave financial institutions bereft of capital for operations and investment. In addition, market volatility and its impact on public confidence can have a significant effect on the broader global economy.

Volatility targeting and risk parity are asset allocation methodologies that are directly impacted by volatility forecasting. Funds that are managed for the insurance industry with a volatility band of 8%–12%—for example—use asset allocation aiming to control the overall fund returns remaining within that range of volatility, as investors with different levels of risk tolerance and time horizons demand differentiated levels of volatility. Maintaining this range of volatility requires that a view be taken on the expected future volatility for the asset classes in the fund. In addition, funds which use risk parity are focused on the allocation of risk rather than the allocation of capital, assuming that each asset class contributes the same degree of volatility to the overall fund. As the volatility of each of these asset classes is not constant, a forecast for the expected volatility for each is required to maintain this type of investment approach.

It is well established that volatility is easier to predict than returns. Volatility possesses a number of stylized facts which make it inherently more forecastable. As such, volatility prediction is one of the most important and, at the same time, more achievable goals for anyone allocating risk and participating in financial markets.

The volatility of asset returns is a measure of how much the return fluctuates around its mean. It can be measured in numerous ways but the most straightforward is historical, observed volatility, which is measured as the standard deviation of asset returns over a particular period of time. When volatility is calculated by reverse-engineering options market prices, it essentially becomes both a market price for and an expectation of uncertainty.

The stochastic or random nature of asset prices and returns necessitates the use of statistics and statistical theory to help describe and predict these market fluctuations. The entire field of financial econometrics is predicated on the integration of the theoretical foundations of economic theory with finance, statistics, probability, and applied mathematics to make inferences about the financial and market relationships critical in the disciplines of asset allocation, risk management, securities regulation, hedging strategies, and derivatives pricing. Volatility is forecastable because of a number of persistent statistical properties.

Volatility Clustering

There is a delay for large or small changes in the absolute value of financial returns to revert back to mean levels. In other words, the magnitude of financial returns have latency—large changes in financial returns tend to be immediately followed by large changes and small changes tend to be immediately followed by small changes. This can lead to volatility clusters over time.

Many studies have found that volatility clustering is likely due to investor inertia, caused by investors' threshold to the incorporation of new information. Except during times of extreme market turmoil, only a fraction of market participants are actually trading in markets at any given point in time. As such, it takes a period of time for these investors to engage in the market and implement their changing views as new information is revealed. The evidence for volatility clustering is shown by the positive serial correlation (correlation of a return series with itself lagged) in the absolute value of returns which eventually decays over a period of observations (Exhibit 1). Volatility clustering can enhance the ability to forecast volatility. This clustering can be shown by plotting a scatter chart of current month versus next month's volatility (Exhibit 2).

Exhibit 1
Autocorrelation of Global Equity Returns

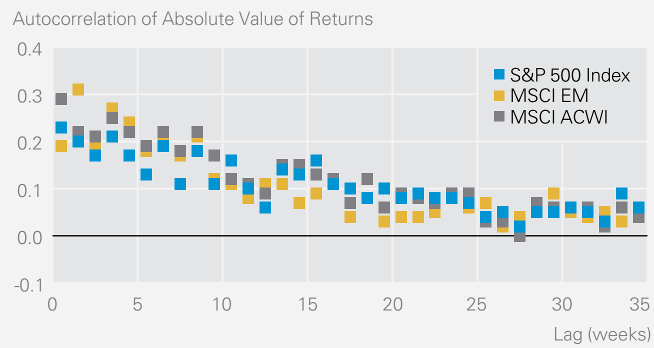

For the period January 1999 to October 2015, weekly returns

The performance quoted represents past performance. Past performance is not a guarantee of future results. This is not intended to represent any product or strategy managed by Lazard. It is not possible to invest directly in an index.

Source: Bloomberg

Exhibit 2
Past Volatility May Be Indicative of Future Volatility…

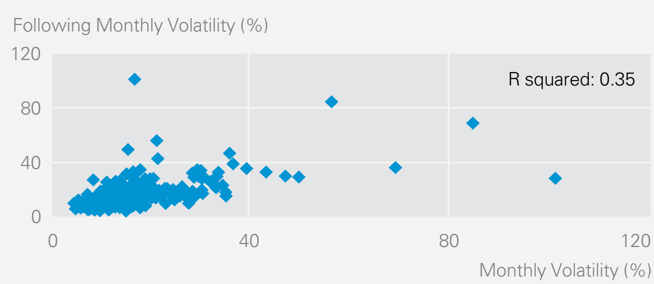

For the period February 1984 to September 2015, monthly returns

Data are based on the S&P 500 Index.

Source: Bloomberg

In contrast, if one plots the current month's return versus the next month there is no linear relationship and the serial correlation of actual returns—not the absolute value of returns—remains insignificant (Exhibit 3).

Exhibit 3
...But Not Indicative of Future Returns

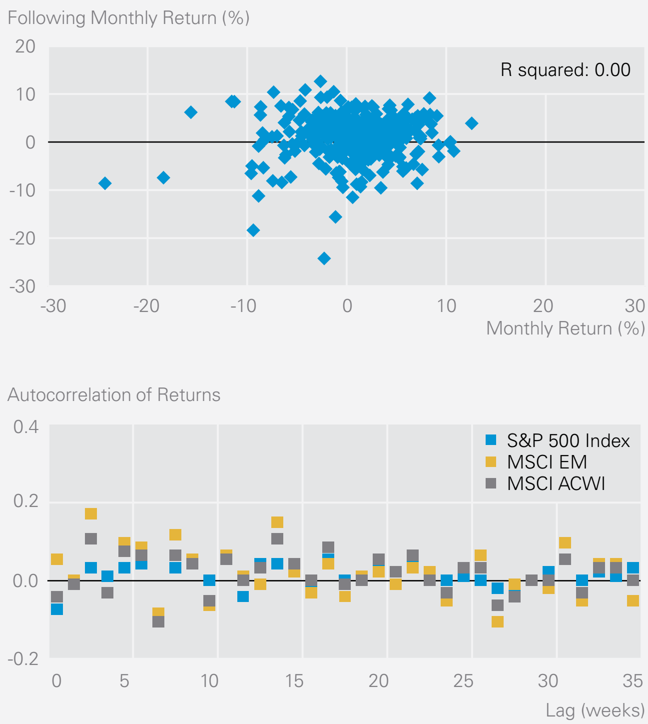

Top chart: for the period February 1984 to September 2015, monthly returns and data are based on the S&P 500 Index.

Bottom chart: for the period January 1999 to October 2015, weekly returns

The performance quoted represents past performance. Past performance is not a guarantee of future results. This is not intended to represent any product or strategy managed by Lazard. It is not possible to invest directly in an index.

Source: Bloomberg

Leverage Effect

The hypothesized leverage effect along with the volatility feedback effect describes the negative and asymmetric relationship between volatility and returns. The mathematical calculation of volatility is indifferent to the direction of the market. However, volatility is negatively correlated to returns. At the same time, negative returns result in larger changes in volatility than positive returns. The beta of the CBOE Volatility Index (VIX) to the S&P 500 Index on negative return days is -3.9 with an r-squared of 0.36 whereas the beta of VIX to the S&P 500 Index on positive return days is -2.8 with an r-squared of 0.23 (Exhibit 4).

The volatility feedback effect suggests that as volatility rises and is priced into the market, there is a commensurate rise in the required return on equity as investors place a higher hurdle rate on returns to achieve their desired risk-adjusted upsides. This leads to an instant decline in stock prices as the volatility immediately reduces the risk-adjusted attractiveness of equities. As stock prices fall, companies become more leveraged as the value of their debt rises relative to the value of their equity. As a result, the stock price becomes more volatile. This effect is more pronounced in well-developed markets that have more analyst coverage.

Mean Reversion

Another stylized property of volatility is that it reverts to the mean over time. The half-life of volatility is measured as the time it takes volatility to move halfway towards its long-term average. Volatility has a half-life of about 15–16 weeks—based on autoregressive models which we will discuss later. With regards to implied volatility, the degree of mean reversion is both asymmetric and accelerated (Exhibit 5). The half-life of VIX mean reversion is about 11 weeks and is considerably less than the half-life for equity returns, which is roughly 15 to 16 weeks (shown by the autocorrelation in Exhibit 1). In addition, VIX mean reversion is far more pronounced when the VIX reaches higher levels than when it dips below its long-term average. So, historically the VIX has dropped with greater frequency and magnitude when elevated than it has increased when at depressed levels. This suggests that

Exhibit 4
There Is a Negative Relationship between Returns and Volatility

Change in VIX on Negative S&P 500 Index Days

Change in VIX on Positive S&P 500 Index Days

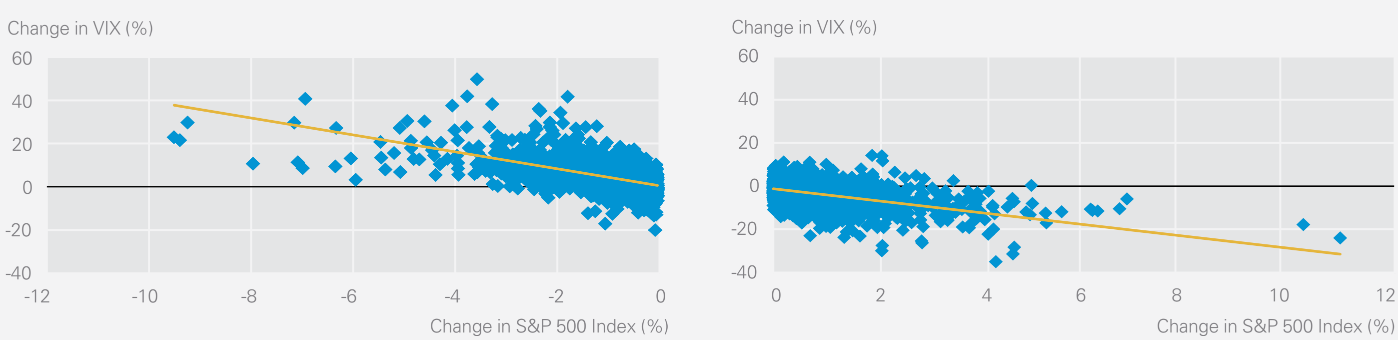

For the period January 1990 to October 2015, daily returns

The performance quoted represents past performance. Past performance is not a guarantee of future results. This is not intended to represent any product or strategy managed by Lazard. It is not possible to invest directly in an index.

Source: Bloomberg

Exhibit 5
VIX Mean Reversion Top/Bottom 10% Converge Towards Long-Term Average

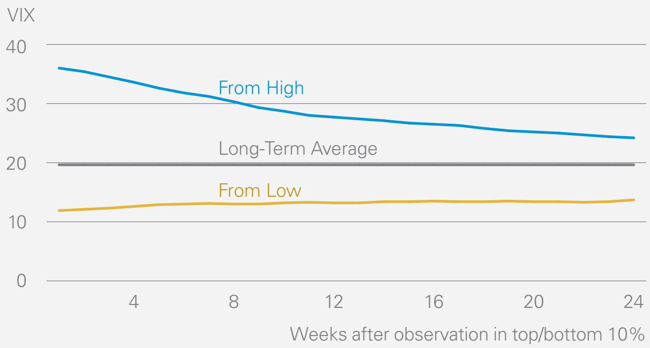

For the period 5 January 1990 to 6 November 2015

We examined the top/bottom 10% of weekly observations. Then we tracked the pattern of the average value from these top and bottom deciles for subsequent weeks.

Source: Haver Analytics

low levels of implied volatility can persist and are much more stable than elevated levels of implied volatility which tend not to persist for as long or at the same magnitude.

The prevailing thesis behind the mean reverting nature of volatility is that in periods of low volatility, investors reduce their expectations and thresholds for volatility and therefore become more sensitive to news flow. This leads to a larger reaction function to new information and higher volatility as a result. Conversely, during periods of elevated volatility, investors will increase their expectations and thresholds for volatility and become less sensitive to new information. This will result in lower levels of volatility in subsequent periods. Finally, this results in the mean reversion with some period of latency as investors gradually adjust their expectations and thresholds to the prevailing levels of market volatility.

Cross Correlation of Volatility

Volatility is correlated across asset classes. Research has shown that the correlation between the volatility of asset classes is stronger than the correlation among the asset returns. Under certain circumstances global bonds and global equities are negatively correlated and in other cases they are positively correlated. The correlation between the US equity and bond market returns has historically fluctuated between -0.6 and +0.8.[2] When asset allocation decisions are made applying certain assumptions about the expected correlation between equities and bonds, changes in these correlations can be a great source of volatility. A more stable relationship can be observed between the volatility of different asset classes than the underlying returns (Exhibit 6). When evaluating different markets—such as bond markets, exchange rates, and equity markets—large movements in one market are often accompanied by large movements in another. The correlation between implied volatility in the equity and bond market fluctuates in a much smaller range around a positive mean.

Methods of Volatility Forecasting

A summary of this section is presented in the box "Summary Characteristics of Volatility Forecasting Methods" (page 7).

HIS–Historical Volatility Models

Historical volatility models are created directly from realized volatility calculated over a specific time frame. These models are the easiest to create and adjust and have strong forecasting performance when compared to more complex models. In general, by adding emphasis to more recent observations by employing a decay factor, these models can incorporate both the persistence and the mean reverting nature of volatility to become a straightforward, robust volatility forecast.

Random Walk

The simplest type of historical volatility model is the random walk model. According to this model, the difference between today's volatility and tomorrow's volatility is just random noise. As such, the best forecast for tomorrow's volatility is today's volatility. This model captures the persistence and changes in near-term levels of volatility. However, it is limited since it does not incorporate the mean reverting nature of volatility. The random walk method is as likely to over-estimate volatility as it is to under-estimate it.

Historical Mean

The historical mean method makes a forecast based on the entire history of volatility. It gives equal weight to all observations. This method captures the mean reverting aspect of volatility but is unable to capture the short-term persistence and changing nature of volatility. It assumes that volatility will immediately revert to its long-run average in the next measurable time period. As a result, it consistently under-estimates future volatility.

Exhibit 6
Correlation between Bond and Equity Markets Is Stronger (and Fluctuates Less) for Volatility than for Asset Returns

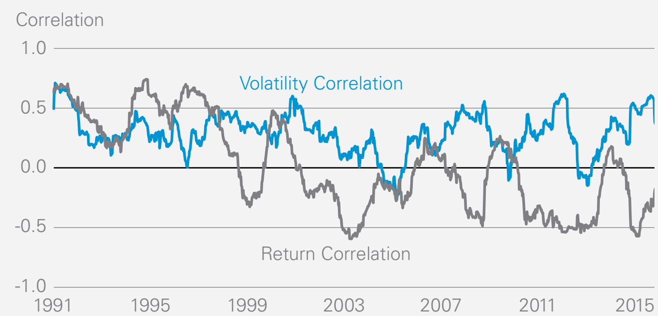

As of January 1990 to October 2015, weekly returns

Volatility correlation: MOVE Index and VIX Index. Return correlation: S&P 500 Index and Barclays US Aggregate Bond Index

Source: Bloomberg

Moving Average (MA)

The next type of historical model is the moving average model. Moving average models estimate a sample mean volatility using a fixed length of time and a fixed weighting scheme. In other words, all observations have the same contribution to the overall volatility calculation. Since it does not include the entire history of volatility, it is more weighted towards recent history and captures more of the recent fluctuations in volatility. By adjusting the weighting scheme, this method can be modified to account for the latency of volatility. The shorter period used for the moving average, the more sensitive to shorter-term fluctuations in volatility and hence the more likely this model is to over-estimate future volatility.

Exponentially Weighted Moving Average (EWMA)

The EWMA model is a historical volatility model which gives more weight to recent observations. This can increase the predictive ability of the model by capturing the tendency of periods of high and low volatility to persist—also known as volatility clustering or the "GARCH[3] effect". This model incorporates the majority of the stylized features of historical return series such as latency of volatility as well as the long-term mean reverting. EWMA helps balance out the over- and under-estimation of volatility.

Discrete Historical Models

Discrete historical models use distinct rolling historical time periods and assign weights to each of these time periods while keeping the weights to each individual observation within a time period static. An example is the aggregate of three moving averages: 22 days, 26 weeks, and 36 months. A weight is applied to each of these three time periods—this weight can be auto-regressed over a specific time period to minimize the volatility estimation error. This is similar to the autoregressive process used in ARMA (autoregressive moving average) and GARCH. However, the difference is that the weights are not adjusted for each observation. Instead, the weights are adjusted for the three groups of observations. In the above case, one could apply different weights to the 22-day period, the 26-week period, and the 36-month period—depending on how much one wants to weight the more recent past versus the medium- and long-term. In this way, some level of specification can be applied but only at the level of each discrete time period. As such, individual observations can only have a limited impact on the specificity of the model. This helps the model maintain a heightened level of exposure to the broader changes in volatility dynamics, while limiting the sensitivity to individual observations and hence tempering the possibility of model over-specification. Including distinct long-, medium-, and short-term time periods, this method captures the clustering nature of volatility as well as its tendency to revert to a long-term mean over time.

Autoregressive Moving Average (ARMA) Models

ARMA models add an autoregression to the moving average used in the typical historical models. ARMA models are moving average models that adjust the weights of the observations to optimize the predictive power over a sample period. These models are well understood and computationally straightforward. There has been significant research and analysis supporting the use of ARMA models for good forecasting performance compared with other models. As discussed, an autoregressive model expresses a time series of returns as a linear function of its past values. The simple regression method is the most straightforward of the ARMA models and takes the historical model one step further by calculating the observation weighting scheme based on a simple regression. So, it essentially adds another lever or variable—the weight applied to the observation—that can be adjusted to better fit the past observations to the subsequent observations. The weighting method becomes dynamic and is no longer pre-specified. In this model, volatility is expressed as a function of its past values along with an error term.

Autoregressive Models

Autoregressive models (AR) are a statistical technique involving a regression of lagged values—the model suggests that past values can help forecast future values of the same variable. Within the model, a time series is the dependent variable and lagged values are the independent variables. For instance, with a five-year monthly time series (60 observations, labeled 1 to 60) the process takes observations 2 through 60 as the independent variable and observations 1 through 59 as the dependent regression variable. This one-lag model is typically referred as AR(1), with the number in parentheses representing the lag. Importantly, for every lag one observation is lost as regression variables need to be the same size.

ARCH Models—Autoregressive Conditional Heteroskedasticity

The ARCH model, was originally developed by Robert Engle in 1982 to measure the dynamics of inflation uncertainty. This finding earned Engle a Nobel Prize and the concept has been applied to many other disciplines, such as medicine.[4] It has been observed that some periods in markets are riskier than others and that these periods are not randomly dispersed across time. In other words, there is a degree of latency or persistence in volatility. Conditional heteroskedasticity refers to the notion that the next period's volatility is conditional on the volatility in the current period as well as to the time varying nature of volatility. In a simple ARCH model, the next period's volatility is only conditional upon the last period's volatility. Consequently, this does not fully capture the persistence of volatility in a period of crisis. In this sense, a simple ARCH model is inferior to the ARMA models previously discussed. To address this shortcoming, ARCH has been extended to GARCH, or Generalized ARCH. The GARCH model is a way of specifying the dependence of the time varying nature of volatility. GARCH incorporates changes in the error term—or fluctuations in volatility—and tracks the persistence of volatility as it fluctuates around its long-term average. Observations are exponentially weighted—more weight is given to more recent observations. GARCH models enable one to incorporate the changes in volatility, the persistence of volatility, as well as account for the non-normality or "fat tails" of financial return series. However, the limitation of GARCH is its inability to respond asymmetrically to falling and rising levels of volatility—an important observable and persistent relationship between volatility and asset returns.

To account for these limitations, many non-linear extensions of GARCH have been developed. Asymmetric GARCH describes models that have been developed to account for the asymmetry of volatility—also known as the leverage effect. This is the observation that volatility increases more for negative market moves than it does for positive market moves. In addition, Multivariate GARCH models have been developed to capture the cross correlation of volatility between different asset classes.

Implied Standard Deviation Models

Implied standard deviation (ISD) models use volatility implied by pricing in the options market. Implied volatility in the options market is widely accepted to be the market's expectation of future volatility. The options market is an extremely broad and liquid market with a wide range of maturities and hence houses a wealth of information on volatility expectations across a multitude of markets over many different time horizons. However, due to the depth and breadth of this market, there are many risk premia that implied volatility incorporates that abstract its connection to the expectation of future volatility—namely tail risk (the pricing of unknown unknowns in the market) as well as the supply and demand dynamics of the options market—factors largely divorced from near-term volatility expectations. In addition, there is significant basis risk between the theoretical value of options and the realized value of options and thus the implied volatility measure (Exhibit 7).

In practice, it is difficult to execute the arbitrage required to ensure that option prices conform to their theoretical value. Critical choices need to be made on what strike price or what maturity to pick, since options with different maturities and strike prices produce vastly different implied volatilities. The most frequently used strike prices are at-the-money or nearest-to-the-money, to limit the amount of tail risk incorporated in the calculation and to help isolate the expected volatility in the market. Over time, the average maturity of the options market has lengthened. This has reduced the reliability of near-term volatility expectations implied by options. Many investors use the VIX Index to aggregate the proximate term structure of options contracts into one measure of near-term volatility expectations. In many ways, this is

Exhibit 7
Difference between the Implied and Realized Value of S&P 500 Index Volatility

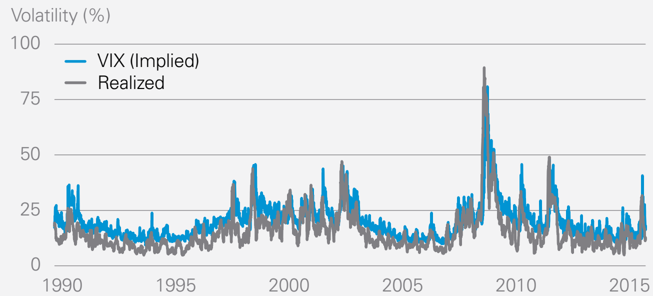

As of January 1990 to November 2015
Source: Bloomberg

tantamount to aggregating the entire yield curve into a single interest rate. As one of the most broadly used hedging instruments in global markets, the VIX is largely driven by tail risk and supply and demand dynamics rather than a pure approximation of near-term volatility.

The ISD method has had mixed results historically. This method has been successful in many studies using single securities.[5] The options written on a particular security contain a specificity impossible to expand to an entire asset class (as is the VIX's case). These are likely a better estimate of future volatility in a particular security as the idiosyncratic risk levels are much more prominent in this case than the general tail risk and supply and demand dynamics that permeate asset class–level options markets. Of note, implied volatility has performed better in interest rate markets than equity markets.

Incorporation of Exogenous Factors/Variables

All of the models described above are univariate—the time series analyzed is the only factor driving the volatility forecast. Clearly, other variables exogenous to the time series can provide an economic or structural explanation for changes in volatility. As such, adjustments can be made to models based on exogenous factors like monetary policy, data releases, recessions, market segment, and interest rates. For example, as US interest rates have an impact on company borrowing costs, higher levels of interest rates can be associated with periods of higher volatility in equity market returns. These methods

Exhibit 8
Forecast Effectiveness

Model	Random Walk	Historical Mean	3M Moving Average	EWMA	ARMA	Discrete Historical	GARCH(1,1)	ISD
Mean Squared Estimation Error	0.0045	0.0080	0.0048	0.0044	0.0043	0.0041	0.0043	0.0057
Mean Over-Estimation Error (%)	4.3	3.9	4.0	4.1	3.7	3.6	3.8	6.1
Mean Under-Estimation Error (%)	-4.8	-8.0	-4.9	-4.8	-4.8	-4.8	-4.9	-5.3
% Periods Over-Estimated	53.0	56.0	55.2	49.1	42.7	52.5	47.1	84.7
Mean Time of Over/Under-estimation (months)	1.53	6.28	2.14	1.57	1.62	1.76	1.61	6.16

For the period February 1988 to October 2015, monthly returns
Forecasted or estimated results do not represent a promise or guarantee of future results and are subject to change
Source: Lazard estimates with data from Bloomberg

Summary Characteristics of Volatility Forecasting Methods

Model	Advantages	Disadvantages	Most Appropriate Asset Classes
Random Walk	• Easy to create • Transparent • Incorporates recent changes in volatility	• Too much focus on near-term realized volatility • Does not capture mean reversion of volatility • Entirely backward looking • Does not capture the asymmetry of volatility	• Does not capture mean reversion characteristic of most asset return series • Limited effectiveness for equities, exchange rates, and interest rates
Historical Mean	• Easy to create • Transparent • Captures long-term mean reversion of volatility	• Only focus is on long-term average volatility • Ignores the near-term persistence and fluctuations in volatility • Does not capture the asymmetry of volatility • Typically over-estimates volatility but the degree of volatility under-estimation is much larger than the degree of over-estimation • Persistent over- and under-estimation of volatility	• Not good for modeling asset classes like equities, interest rates, and exchange rates which exhibit changes and clustering in volatility
Moving Average	• Easy to create • Transparent • Captures long-term mean reversion of volatility • Captures some near- term persistence and fluctuations in volatility	• Weights all observations equally so near-term fluctuations in volatility are not captured sufficiently • Does not capture the asymmetry of volatility	• Equity indices
Exponentially Weighted Moving Average (EWMA)	• Easy to create • Transparent • Captures long-term mean reversion of volatility • Captures near-term persistence and fluctuations in volatility • Decay factor can be adjusted to time series based upon the differential half-life of a specific return series	• Does not capture the asymmetry of volatility • A decay factor that is optimized to a particular time period can create over-specification and result in reduced predictability in out-of-sample time periods	• Equities • Equity indices • Interest rates • Exchange rates
Autoregressive Moving Average (ARMA)	• Captures long-term mean reversion of volatility • Captures near-term persistence and fluctuations in volatility • Autoregressive—the weight applied to the observation can be adjusted to better fit the past observations to the subsequent observations	• Does not capture the asymmetry of volatility • Risk of over-specificity	• Equities • Equity indices • Interest rates • Exchange rates
Discrete Historical	• Captures long-term mean reversion of volatility • Captures near-term persistence and fluctuations in volatility • Very applicable to uses in asset allocation	• Does not capture the asymmetry of volatility	• Equities • Equity indices • Interest rates • Exchange rates
General Autoregressive Conditional Heteroskedasticity (GARCH)	• Captures long-term mean reversion of volatility • Captures near-term persistence and fluctuations in volatility • Autoregressive—the weight applied to the observation can be adjusted to better fit the past observations to the subsequent observations • Can be modified to account for the asymmetry of volatility • Can be multivariate to capture the cross-correlation of volatility across asset classes	• Risk of over-specification to a particular sample period reducing forecasting abilities in out-of-sample time periods	• Equities • Equity indices • Asymmetric; GARCH is uniquely able to capture the "leverage effect" in equity market returns • Exchange rates
Implied Standard Deviation Models (ISD)	• Implied volatility contains useful information about the market's expectations for future volatility • Very useful for predicting the expected volatility of individual securities • Captures long-term mean reversion of volatility • Captures near-term persistence and fluctuations in volatility	• There are many risk premia incorporated in implied volatility measures in addition to proximate volatility expectations • Consistent and persistent over-estimation of volatility • There is significant basis between the theoretical value of options and the realized value of options and thus the implied volatility measure • Difficult to aggregate the volatility implied by the options market into one measure of expected volatility for an entire asset class	• Individual equities • Individual currency exchange rates • Provides valuable information on the future volatility of equity indices; however, model-based historical forecasts typically have better predictive ability

largely involve the use of ARCH models augmented to include various uncorrelated external risk drivers that are likely to affect volatility. These risk drivers include economic factors related to credit, inflation, interest rates, economic growth, liquidity, and currencies. Exogenous factors can have a significant and non-linear impact on realized volatility. Over the short term, the predictive power from including exogenous factors typically matches that of GARCH models. However, over the longer term, incorporating these exogenous factors in a non-linear fashion adds to the predictive power of the model.

Results—Forecast Effectiveness

Exhibit 8 displays the predictive results from methods discussed. The discrete historical, ARMA, GARCH (1,1) and EWMA are the most accurate forecasting methods based on the lowest mean squared estimation error (i.e., the squared difference between forecasted and realized values) while the historical mean method had the highest estimation error. The common factor contributing to the lower estimation errors of the discrete historical, ARMA, and GARCH (1,1) models is the additional layer of specification that an autoregressive weighting scheme provides. The ISD method was the most likely to over-estimate future volatility (84.7% of the time) by the highest magnitude (6.1%) and with a high level of persistence (6.2 months of average over- and under-estimation). This is largely expected due to the additional risk premia embedded in measures of implied volatility. ARMA was most likely to underestimate future volatility but the magnitude and persistence of that underestimation was relatively low. These results indicate a clear tradeoff between the additional forecast accuracy that an autoregressive weighting scheme provides and the risk of over-specification to a particular sample period.

In addition, we tested the models against the realized volatility of the S&P 500 Index during the heights of the global financial crisis in 2008 and the euro zone sovereign debt crisis in 2011 (Exhibit 9). The quickest model to react to rising levels of volatility in a crisis is the random walk model—which only incorporates recent changes in volatility. The slowest to respond is the historical mean model—which only incorporates the long-run average of volatility. The rest of the models' volatility projections during these crisis periods fall in between. Interestingly, while consistently over-estimating volatility over long time periods, the implied volatility model actually underestimated volatility during the global financial crisis—achieving a peak volatility of 60% using monthly VIX levels and 80% using daily VIX levels, while the S&P 500 Index actually reached a realized volatility level of 83%.

Exhibit 9
Volatility Forecasting Methods, 2008–2012

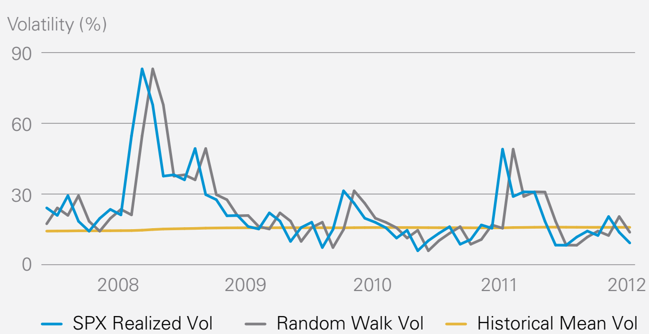

For the period January 2008 to August 2012

Forecasted or estimated results do not represent a promise or guarantee of future results and are subject to change.

Source: Lazard estimates with data from Bloomberg.

Conclusion

Uncertainty is the basis of all financial models and volatility is accepted as the realization of that uncertainty. Several characteristics of financial return series make volatility inherently predictable. However, due to the stochastic nature of volatility, there is no amount of past data that we can plug into a model that will fully capture the current or future behavior of volatility. At best, a model can approximate the behavior of volatility during the sample period analyzed—and forecast accuracy is determined by testing out of sample.

Research has shown that the predictive ability of a model depends largely on the asset class and the frequency of the observations.[6] Some models are better at predicting equity market volatility while others are better at predicting interest rates or currencies. Each successive level of complexity added to these volatility predictive models increases the degree of specificity. From simple historical models, to ARMA models to the many extensions of GARCH, each successive model aims to capture another nuance of the return series over a specific sample period. While this increases the fit of the model to the sample data during the period tested, research has shown that this does not necessarily enhance the model's predictive ability outside of sample periods. Complexity can also engender a false sense of overconfidence in the model's predictions.

Consequently, simpler, broader models which are able to capture the more general features of volatility and financial returns are likely to provide more robust and transparent predictive abilities over longer, out-of-sample time horizons. In particular, these models will incorporate the primary characteristics of historical return series—volatility clustering, the leverage effect, and mean reversion. We prefer a historical method of forecasting volatility which incorporates long-, medium- and short-term realized volatility. Including the historical returns over 3-year, 6-month, and 1-month time periods and including them as discrete allocations to a volatility forecast, and the most prominent stylized facts about historical financial return series—the clustering and long-term mean reversion of volatility. Regardless of the method chosen, volatility forecasting is one of the most profound methodologies in financial economics, as it is a key tool for asset allocation in general, and specifically for investors who implement volatility targeting and risk parity strategies.

This content represents the views of the author(s), and its conclusions may vary from those held elsewhere within Lazard Asset Management. Lazard is committed to giving our investment professionals the autonomy to develop their own investment views, which are informed by a robust exchange of ideas throughout the firm.

References

Blair, Bevan, Ser-Huang Poon, and Stephen J. Taylor. 2001. "Forecasting S&P 100 Volatility: The Incremental Information Content of Implied Volatilities and High Frequency Index Returns." *Journal of Econometrics,* vol. 105 pp. 5–26.

Campbell, John Y., Andrew Lo, and Craig MacKinlay. 1997. *The Econometrics of Financial Markets.* Princeton, NJ: Princeton University Press.

Christensen, Bent and Nagpurnanad Prabhala. 1998. "The Relation between Implied and Realized Volatility." *Journal of Financial Economics,* 50:2, pp.125–50.

Fleming, Jeff, Barbara Ostdiek, and Robert E. Whaley. 1995. "Predicting Stock Market Volatility." *Journal of Futures Markets,* 15:3, pp. 265–302.

Fleming, Jeff. 1998. "The quality of market volatility forecasts implied by S&P 100 Index option prices." *Journal of Empirical Finance,* 5, 4, pp. 317–345.

Engle, Robert F. 1982. "Autoregressive Conditional Heteroskedasticity with Estimates of the Variance of United Kingdom Inflation." *Econometrica,* 50:4, pp. 987–1007.

Martin-Guerrero, J.D., et al. 2003. "Dosage individualization of erythropoietin using a profile-dependent support vector regression." *IEEE Transactions on Biomedical Engineering,* Vol. 50, issue 10, pp. 1136–1142.

Poon, Ser-Huang and Clive W. J. Granger. 2003. "Forecasting Volatility in Financial Markets: A Review." *Journal of Economic Literature,* Vol. 41, no. 2, pp. 478–539.

Notes

1 Campbell, Lo, and MacKinlay (1997)

2 Data are for the period 10 January 1990 to 28 October 2015, based on the S&P 500 Index and the Barclays US Aggregate Bond Index, weekly frequency.

3 Generalized Autoregressive Conditional Heteroskedasticity

4 Dynamic volatility models are used to find the optimal EOP dosage for anemia patients (Martin-Guerrero et al. 2003).

5 See for example: Fleming, Ostdiek, and Whaley (1995); Christensen and Prabhala (1998); Fleming (1998); Blair, Poon, and Taylor (2001).

6 Poon and Granger (2003)

Important Information

Published on 28 December 2015.