

CONTRIBUTORS

Tim Edwards, PhD
Senior Director
Index Investment Strategy
tim.edwards@spglobal.com

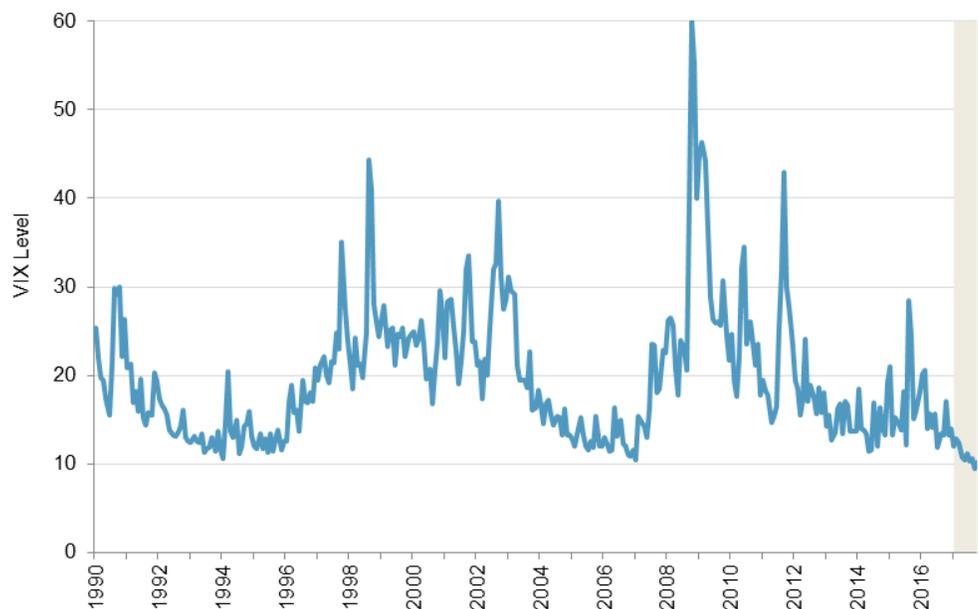
Hamish Preston
Senior Associate
Index Investment Strategy
hamish.preston@spglobal.com

Reading VIX[®]: Does VIX Predict Future Volatility?

EXECUTIVE SUMMARY

The CBOE Volatility Index, otherwise known as “VIX,” is a measure of anticipated movements in the [S&P 500[®]](#), derived from the current traded prices of S&P 500 options.¹ Known as Wall Street’s “fear gauge,” VIX is followed by a multitude of market participants, while its levels and trends have become part of the common language of market commentary. Exhibit 1 shows VIX levels thus far in 2017 (shaded section) in their historical context.

Exhibit 1: Historical VIX Levels



Source: S&P Dow Jones Indices LLC and CBOE. Data from January 1990 to October, 2017. VIX levels are as of the last trading day of each month. Past performance is no guarantee of future results. Chart is provided for illustrative purposes.

Like other indices, the level of VIX is determined by the price of a basket of tradable constituents—in particular, a basket of options that expire in the next month or so. The profit or loss that option buyers and sellers achieve during the life of those options will, among other factors, depend on how significantly the actual volatility of the S&P 500 differs from the volatility “implied” by VIX at the start of the period. If VIX is too low, market

¹ For details on how VIX is calculated, please see <https://www.cboe.com/micro/vix/vixwhite.pdf>.

The purpose of this paper is to provide—without requiring a prior knowledge of the sophisticated mathematics involved in option pricing—a guide to interpreting what is, and what is not, indicated by VIX

participants might profit by buying options, and if it is too high, market participants might profit by selling options. Thus, VIX may be thought of as a crowd-sourced estimate for the anticipated volatility of the S&P 500. In the same way that one cannot invest directly in an interest rate or a dividend, even though one can speculate on their future value, it is impossible to invest directly in VIX, and the meaning of a given VIX level is frequently misunderstood.

The purpose of this paper is to provide—without requiring a prior knowledge of the sophisticated mathematics involved in option pricing—a guide to interpreting what is, and what is not, indicated by VIX. Specifically, we shall:

- Explain how VIX may have greater significance when viewed *relative* to the recent historical level of S&P 500 volatility;
- Describe the dynamics of mean reversion and the premium underlying the typical relationship between VIX and recent volatility levels;
- Derive an estimate for the level of “expected” VIX at any point in time, based on historical norms, and explain how the difference between the expected and actual VIX levels may be interpreted; and finally
- Examine the power of VIX and other related indicators in making predictions, in the U.S. and across various global markets.

VIX retains meaningful predictive aspects, and provides useful indications of sentiment, but they must be carefully teased from the data.

Our ultimate aim is to provide ways to interpret VIX in a manner that teases out a better prediction for the absolute levels of future volatility and a more meaningful gauge of market sentiment.

INTRODUCTION

“For me context is the key—from that comes the understanding of everything.”

- Kenneth Noland

In recent months, many column inches have been devoted to the observation—and speculation upon the meaning—of unusually low VIX readings. Some have been perplexed by the apparent contradiction between a low VIX and their perceptions of a heightened risk environment. Many have begun to question if VIX has become complacent or lost its relevance.²

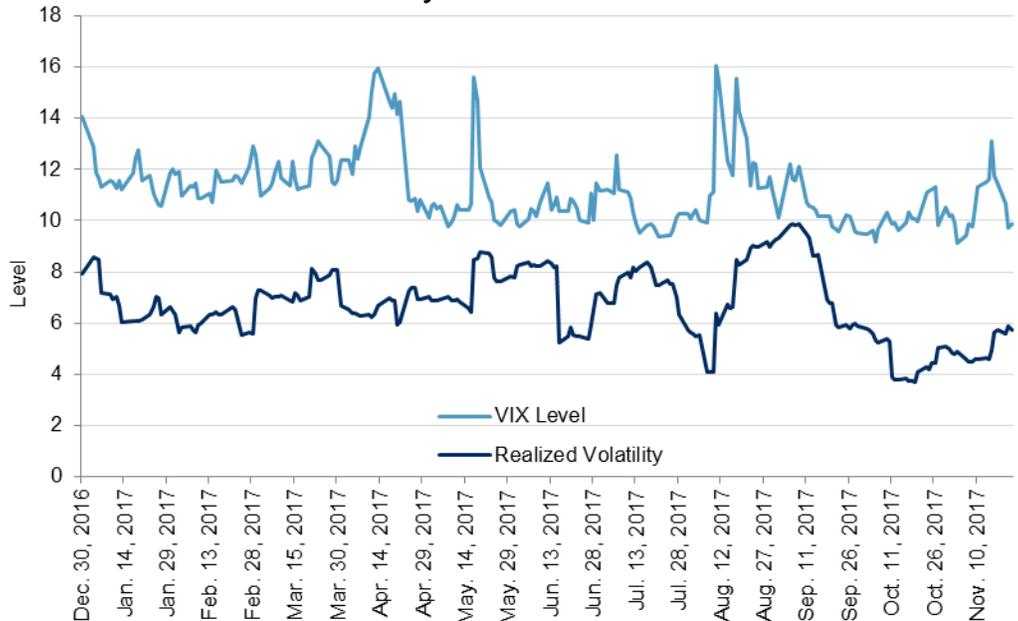
Was 2017’s low VIX (shown in Exhibit 1) indicative of undue complacency? It appears not. Exhibit 2 shows VIX and the corresponding level of S&P

² For example, the *Financial Times* reported “[Worries over complacency as VIX slips to year low](#)” (Dec. 21, 2016) and “[The fearless market ignores perils ahead](#)” (April 18, 2016).

500 realized volatility³ between Dec. 31, 2016, and Nov. 22, 2017. Since the VIX level at any point may be naively interpreted as a prediction for the annualized level of realized volatility over the next 30 days, **a better question might be why VIX has been so high**—averaging more than four points higher than actual realized volatility.

Noting that the VIX level at any point may be naively interpreted as a prediction for the annualized level of realized volatility observed 30 days later, a better question might be why VIX has been so high, averaging over four points higher than realized volatility, as opposed to so low?

Exhibit 2: VIX and Recent Volatility



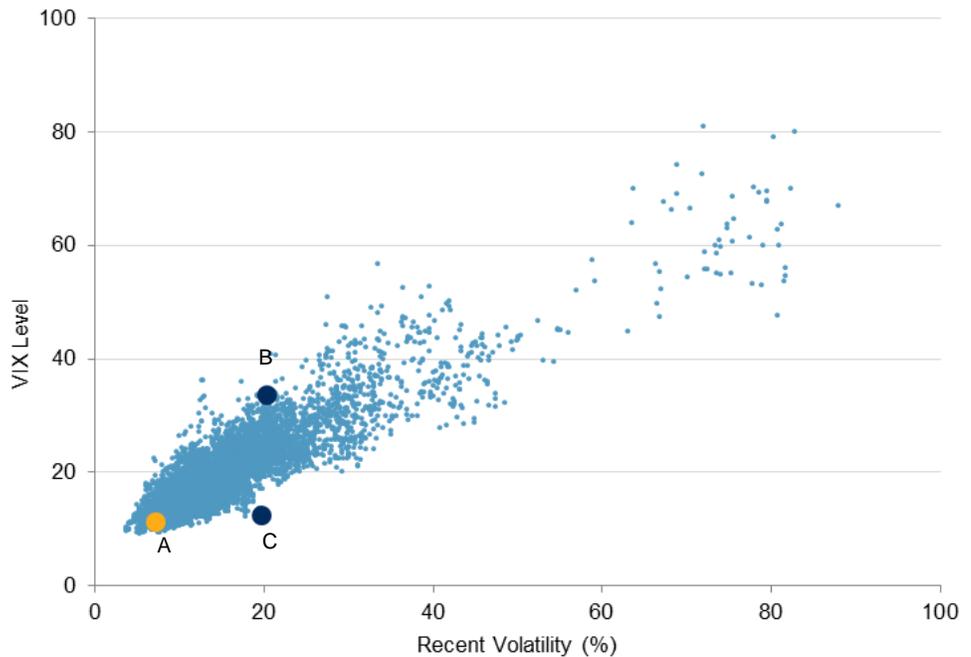
Source: S&P Dow Jones Indices LLC and CBOE. Data from Dec. 31, 2016, to Nov. 22, 2017. VIX levels and S&P 500 realized volatility levels, as calculated using closing price levels, are given for each trading day. Past performance is no guarantee of future results. Chart is provided for illustrative purposes.

Hereafter, as in Exhibit 2, we shall refer to trailing 30-calendar day realized volatility as the “**recent volatility**” corresponding to that point in time. Later sections will examine the extent to which VIX overestimates future volatility on a systematic basis and what drives this overestimation; for now, **we emphasize the usefulness of comparing VIX to realized volatility in judging whether VIX is “high” or not.**

To provide a longer-term perspective, Exhibit 3 shows the historical relationship between VIX and recent volatility over the past 27 years. Point A shows the average levels of VIX and recent volatility observed over the first 10 months in 2017.

³ Specifically, annualized volatility levels in the S&P 500 as measured via daily closing log price changes over the prior 30 calendar days; see Appendix A for more details.

Exhibit 3: VIX and Recent Volatility



Source: S&P Dow Jones Indices LLC and CBOE. Data from Jan. 1, 1990, to Oct. 31, 2017. Chart is based on VIX levels and their corresponding S&P 500 recent volatility levels on each trading day. Past performance is no guarantee of future results. Chart is provided for illustrative purposes.

Exhibit 3 shows that the VIX level is clearly associated with the level of recent volatility. However, there is a significant degree of variation; for example, a recent volatility level of around 20% has, at one point in time, accompanied a VIX level of 34 (point B, when we might say that VIX was relatively “high”), and at another point in time, a VIX level of 12 (point C, when we might say VIX was relatively “low”). The following sections aim to formalize such observations, providing a test for whether VIX is high or low, and a mechanism to interpret such differences.

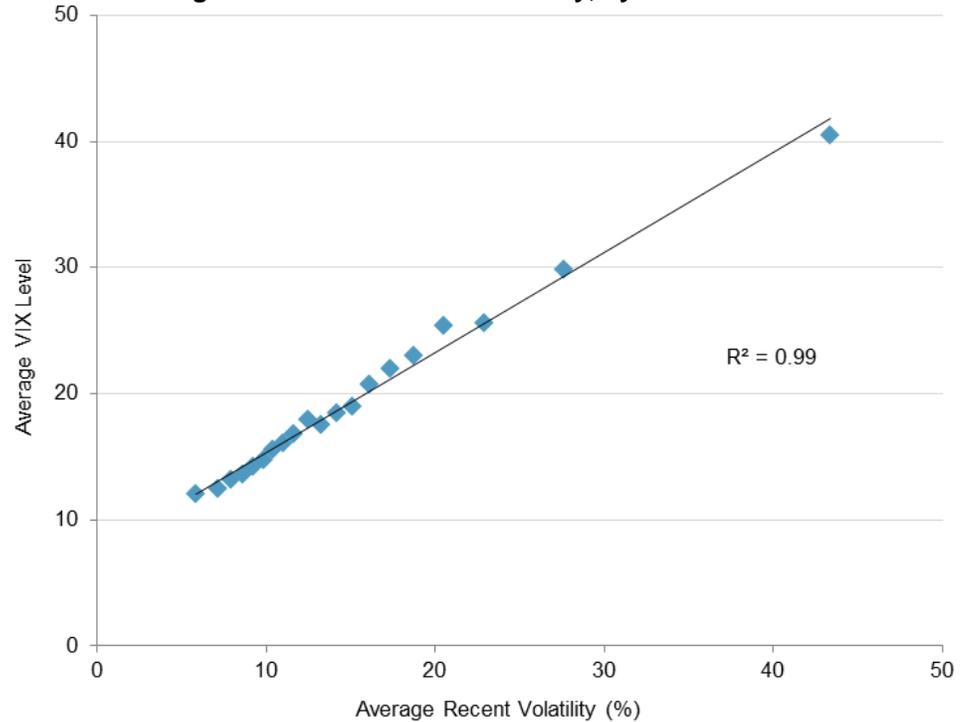
REALIZED VOLATILITY AND VIX

Although Exhibit 3 might suggest that VIX varies in a roughly “straight line” fashion with recent volatility, it is entirely possible that the variation in the data masks a more nuanced relationship. In other words, there is enough visible noise in the data to suppose a degree of curvature is possible in their dependence. A simple way to assess the suitability of a *linear* relationship is to examine the relationship between local averages.

Specifically, we grouped the data of Exhibit 3 into 20 “buckets”: those days when recent volatility was in the bottom 5% of all such observations, those days when recent volatility was in the bottom 6%-10% of all such observations, and so on up until the final category containing the days when recent volatility was in the highest 5% of all observations. The average recent volatility level and the average VIX level in each of the 20 buckets are plotted in Exhibit 4.

A simple way to assess the suitability of a linear relationship is to examine the relationship between the local averages.

Exhibit 4: Averages of VIX and Recent Volatility, by Percentile



A linear relationship between VIX and recent volatility seems feasible.

Source: S&P Dow Jones Indices LLC and CBOE. Data from Jan. 1, 1990, to Oct. 31, 2017. Chart is based on VIX levels and their corresponding S&P 500 recent volatility levels on each trading day. Past performance is no guarantee of future results. Chart is provided for illustrative purposes.

Exhibit 4 makes it clear that the assumption of a linear (or at least *near-linear*) relationship between recent volatility and VIX levels is quite reasonable. In other words, Exhibit 4 suggests that, in estimating whether the current VIX is “high” or “low,” it may be productive to **compare current VIX to a linear interpolation of the historical VIX values at similar levels of recent volatility.**

Even after accounting for the historically low realized volatility environment, VIX was lower than we would have expected in 2017. A *combination* of both structural and sentimental causes for “low” VIX levels in 2017 is quite possible.

For example, the average recent volatility level during the first ten months of 2017 was 7%, which would lead us (on a historical basis) to anticipate an average VIX level of approximately 12% over the same period. Instead, the average VIX level between January and October 2017 was 11% (point A of Exhibit 3). As such, VIX **was somewhat lower than we would have expected**, even after accounting for the historically low realized volatility environment.

There are two non-exclusive interpretations of such a low VIX. On the one hand, the lower-than-expected VIX could be taken to indicate that there have been more (or more enthusiastic) market participants selling options (or volatility) in 2017 than was historically typical. Alternatively, there may not have been a significant change in the demand and supply for options, but instead general investor sentiment may have incorporated a prediction for continued lows in recent volatility (quite correctly, as it turned out). Of course, a *combination* of both structural and sentimental causes for the low VIX levels in 2017 is quite possible, although as we shall show in later

sections, a lower-than-expected VIX has historically proved an occasionally useful indicator of future declines, or lows, in recent volatility.

DECOMPOSING VIX

“Divide each difficulty into as many parts as is feasible and necessary to resolve it.”

- Rene Descartes.

Viewing VIX levels in the context of recent volatility provides us with an assessment of whether VIX is different from “usual.” **But how might we explain differences between observed and expected VIX levels, and what (if anything) does a high or low VIX imply about the future level of realized volatility?**

Answering this question requires us to identify some of the components underlying the relationship between recent volatility and VIX. This decomposition involves several steps, and the result will be a partition of any VIX reading into a sum of the following four distinct components, only the last of which is based on “forward looking” inputs:

We decompose the VIX level into the sum of four distinct components, only one of which encodes forward-looking aspects.

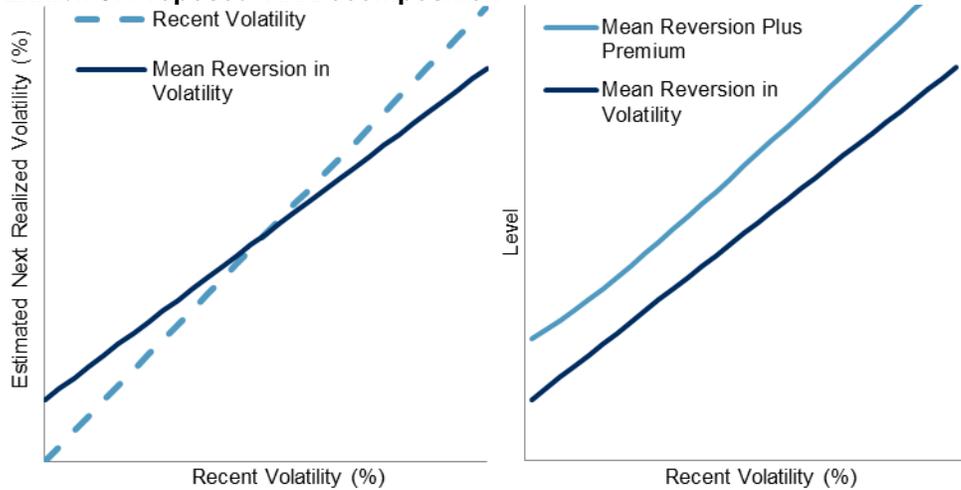
1. The recent volatility environment, *plus*
2. An anticipated (positive or negative) change in recent volatility, predicated on the assumption that volatility reverts at a certain speed towards its long-term mean, *plus*
3. An always-positive “volatility premium” that varies in a predictable manner with recent volatility, although in a not-quite-linear fashion, *plus*
4. A positive or negative component that we shall call the “difference to model,” adjusting for the market’s expectations regarding the magnitude, impact, and frequency of market-moving events in the next 30 days, as well any adjustments to the typical volatility premium.

Exhibit 5 illustrates the first three stages of our proposed VIX decomposition. We begin with the assumption that the first component of VIX is recent volatility—this is the light blue dotted line in the left-hand chart. The second stage is an increase or decrease according to whether recent volatility is above or below the mean level to which it reverts: the dark blue line on the left-hand chart. The right-hand side shows the third stage: the addition of the nearly linear volatility premium, which is positive for all levels of recent volatility in the case of the S&P 500.⁴

This decomposition of VIX will require two main stages, as indicated in Exhibit 5. We begin with mean reversion.

⁴ Please see Appendix, section D for examples of a negative volatility premium.

Exhibit 5: Proposed VIX Decomposition



Source: S&P Dow Jones Indices LLC. Charts are provided for illustrative purposes.

A key characteristic of volatility (realized, or implied) is that it has shown a tendency to revert to its mean.

MEAN REVERSION IN VOLATILITY

“The most important of these rules is the first one: the eternal law of reversion to the mean in financial markets.”

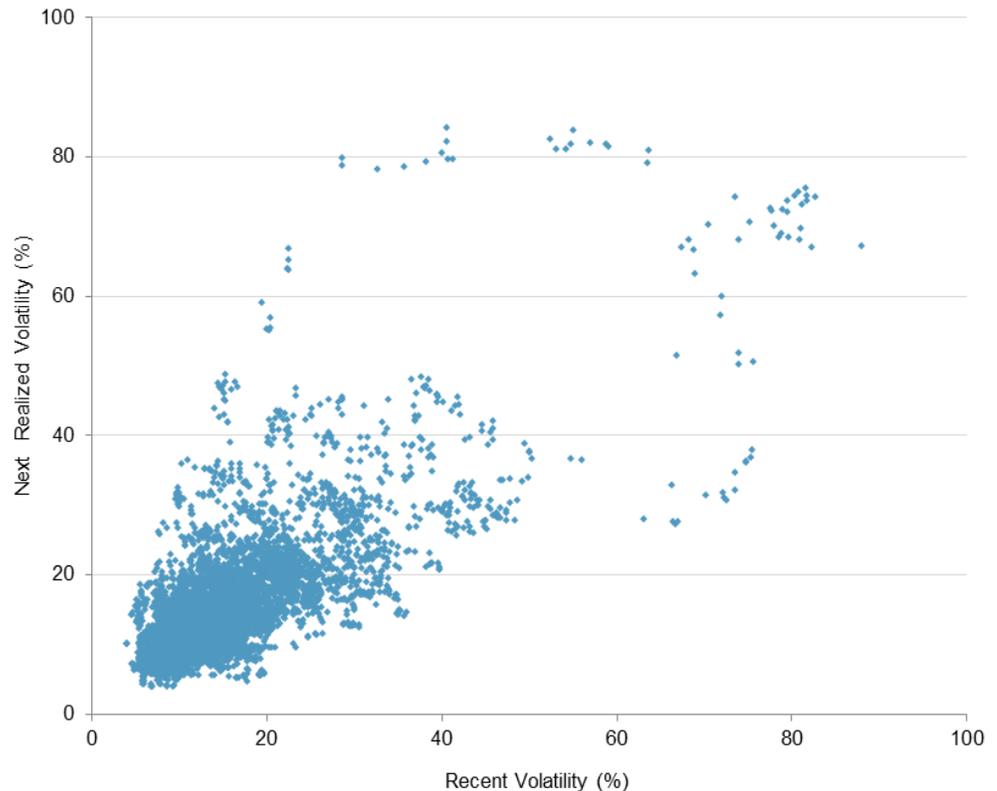
- John Bogle

A key characteristic of volatility (realized or implied) is that it shows a tendency to revert to its mean. This observation is not particularly novel, although it does have a celebrated history. There is an overwhelming bank of evidence to support the mean reversion of volatility across different markets, and the pioneers of research in this area were awarded a Nobel Prize in part for incorporating their findings in the form of volatility predictions and simulations.⁵

In order to demonstrate mean reversion in the volatility of the S&P 500, Exhibit 6 shows the historical relationship between recent volatility at each point in time, and the level of recent volatility observed one month later. Specifically, recent volatility was calculated on each trading day between Jan. 2, 1990, and Sep. 29, 2017, and it was then compared to recent volatility on the last trading day within the following 30 calendar days. We call this “**Next Realized Volatility.**”

⁵ Evidence for mean reversion in several major indices that have associated VIX-like volatility indicators is provided in Appendix C. For an overview of stylized facts regarding asset price volatility, see Engle, Robert F. and Andrew J. Patton, “[What good is a volatility model.](#)” *Quantitative finance* 1, no. 2 (2001): 237-245.

Exhibit 6: Recent Volatility and Next Realized Volatility in the S&P 500

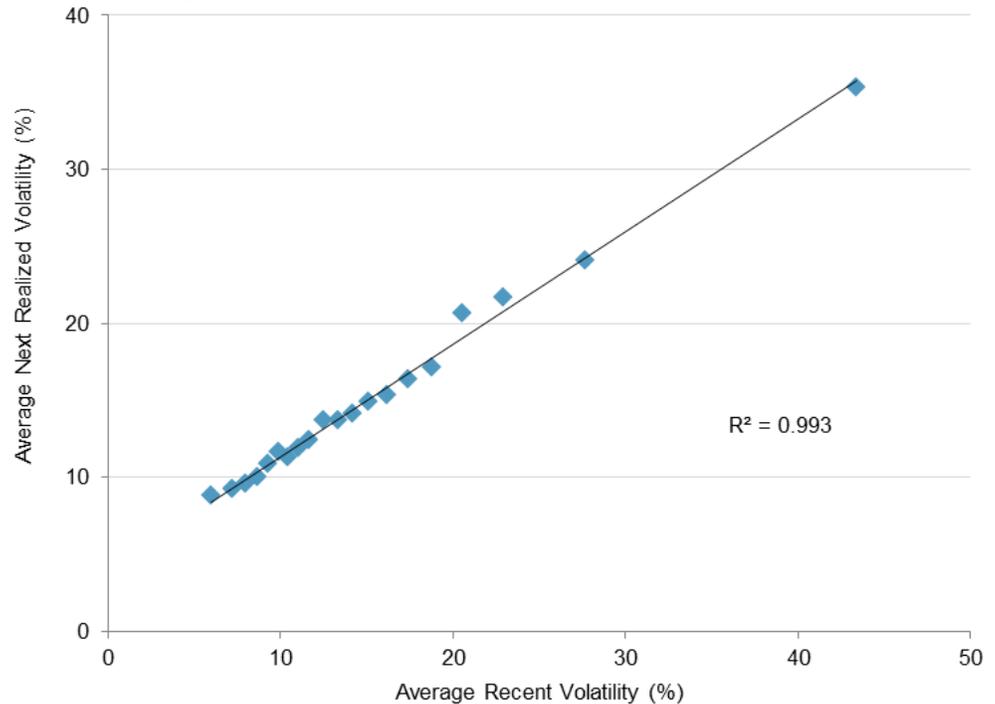


Indeed, there is considerable dispersion in the data; a recent volatility level of 20% has corresponded to next realized volatility levels as low as 6% and as high as 56%, for example.

Source: S&P Dow Jones Indices LLC. Data from Jan. 1, 1990 to Oct. 29, 2017. Chart is based on S&P 500 recent volatility levels and their corresponding next realized volatility levels on each trading day. Past performance is no guarantee of future results. Chart is provided for illustrative purposes.

Beyond an apparently positive relationship around the leading diagonal (which corresponds to the observation that the levels of recent volatility observed one month apart are often similar), the functional form relating the two variables in Exhibit 6 is less clear than it was in Exhibit 3. There is considerably greater dispersion in the data; a recent volatility level of around 20% has corresponded to a next realized volatility level as low as 6% and as high as 56%, for example. With greater scope for a non-linear relationship to exist, we once again compared local averages of the variables. In exactly the same way that Exhibit 4 re-interpreted the data shown in Exhibit 3, Exhibit 7 separates the data of Exhibit 6 into 20 equally sized buckets based on the percentile ranges of recent volatility, and plots the local averages in each of the buckets.

Exhibit 7: Average Recent Volatility and Average Next Realized Volatility in the S&P 500, by Percentile



Source: S&P Dow Jones Indices LLC. Data from Jan. 1, 1990 to Oct. 29, 2017. Chart is based on VIX levels and their corresponding S&P 500 recent volatility levels on each trading day. Past performance is no guarantee of future results. Chart is provided for illustrative purposes.

Exhibit 7 demonstrates that **there was a positive and near-perfect linear relationship, on average, between recent volatility and next realized volatility**. This fact is remarkably convenient; a good fit to a linear relationship of the form $y = ax + b$, where x is recent volatility and y is next realized volatility, means that we can define variables S and M representing the “speed” of reversion and its eventual destination (the “mean”), respectively,⁶ so that line of regression may be written as:

$$\text{Next Realized Volatility} = \text{Recent Volatility} + S * (M - \text{Recent Volatility})$$

Since $(M - \text{Recent Volatility})$ is recognizable as the “distance” between recent volatility and M , the equation may be interpreted as implying that next realized volatility will—on average—match what would be anticipated should recent volatility move a fixed proportion (S) of the distance from its present level, to its mean (M).

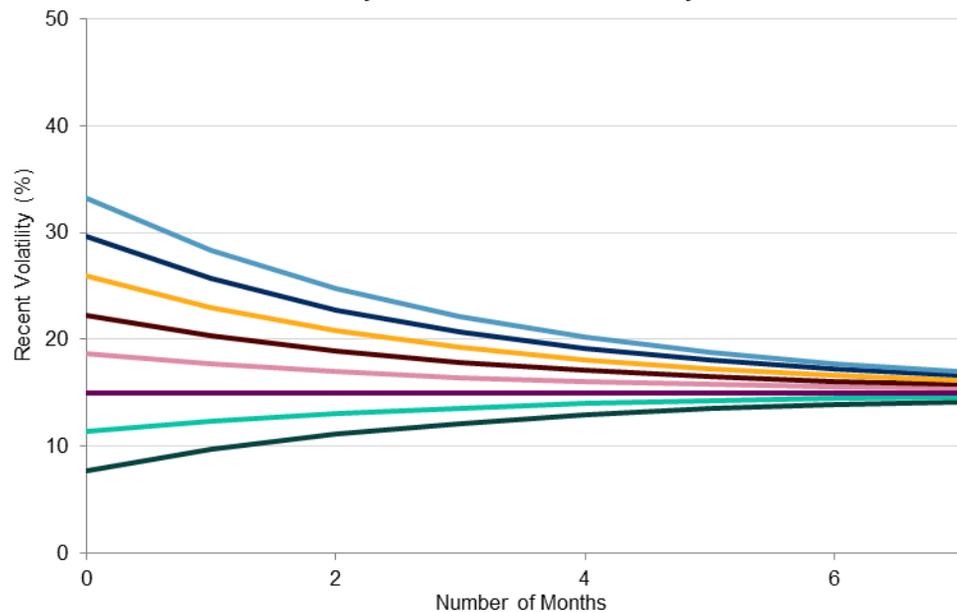
Calibrating to the data of Exhibit 6, $M = 15\%$ and $S = 27\%$ for the S&P 500, meaning that **(on average) S&P 500 recent volatility reverts towards a mean of around 15%, and may be expected to move 27% of the way towards M over a one-month period.**

There was a positive and near-perfect linear relationship, on average, between recent volatility and next realized volatility.

⁶ Formally, this may be achieved by setting $S = 1 - a$ and $M = \frac{b}{1-a}$. In the case of Exhibit 7, $a = 0.73$ and $b = 4.04$ and so $M = 15\%$ and $S = 27\%$, to the nearest percent.

Exhibit 8 plots the typical volatility path under mean reversion of this form. Specifically, recent volatility was given eight values at the start, ranging from 5% to 40% and increasing in increments of 5%. We then applied the mean reversion dynamic to each of these volatility levels over the next seven months such that each of the lines shows the evolution of recent volatility, assuming mean reversion is the only consideration.

Exhibit 8: Mean Reversion Dynamic in Recent Volatility

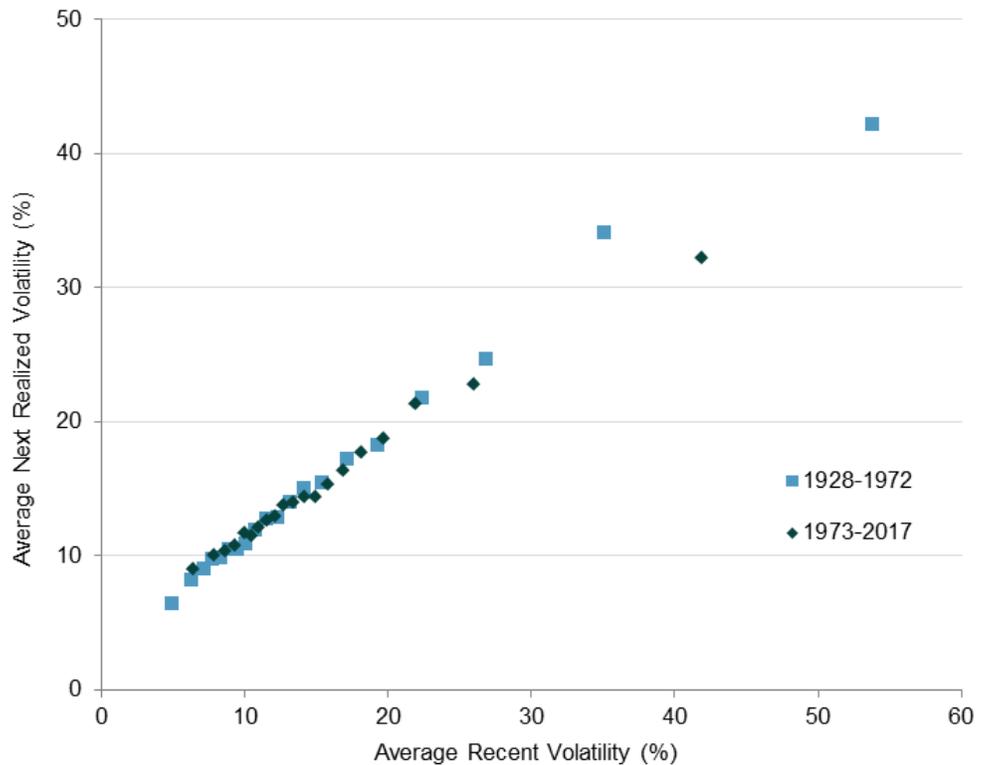


Source: S&P Dow Jones Indices LLC. Chart is provided for illustrative purposes. Chart assumes a constant $M = 15\%$ and $S = 27\%$ over the eight-month period.

Exhibit 8 shows that changes in volatility from month to month are larger when recent volatility is substantially different from $M = 15\%$. It also demonstrates that all of the lines converge toward 15% as time increases. Such a result is not caused by 15% being in any way *special*; this value for M was based on the historical levels of S&P 500 recent volatility and their evolution. It is not beyond the bounds of feasibility that M (otherwise recognized as long-term average U.S. equity market volatility) might change over time; changes to sector weightings in the S&P 500, trading volumes, and regulations all have the potential to impact both the speed and the destination of mean reversion.

Nonetheless, Exhibit 9 suggests that the relationship between recent volatility and next realized volatility for distinct periods in the S&P 500 has a degree of historical stability—at least if measured over a sufficiently long period. Exhibit 9 is calculated in the same way as for Exhibit 7, with the difference being that the light blue dots are based on S&P 500 closing prices from Jan. 27, 1928, to Dec. 29, 1972 and the dark blue dots are based on data from Jan. 2, 1973, to Sept. 30, 2017.

Exhibit 9: Average Recent Volatility and Average Next Realized Volatility, by Percentile



There is an apparent stability in the dynamics of mean reversion over long time horizons, and we suppose that a historically typical degree of mean reversion might therefore be incorporated into volatility expectations.

Source: S&P Dow Jones Indices LLC. Data from Jan. 27, 1928, to Oct. 29, 2017. Chart is based on S&P 500 recent volatility levels and their corresponding S&P 500 next realized volatility levels. These levels are separated into two distinct periods: Jan. 27, 1928, to Dec. 29, 1972, and Jan. 2, 1973, to Oct. 29, 2017. Past performance is no guarantee of future results. Chart is provided for illustrative purposes and reflects hypothetical historical performance. Please see the Performance Disclosure at the end of this document for more information regarding the inherent limitations associated with back-tested performance.

Over the two periods, half a century apart, the regression lines are encouragingly congruent; $M = 16\%$ and $S = 24\%$ for the earlier period, in comparison to $M = 15\%$ and $S = 32\%$ for the later period. Thus, assuming there is a degree of stability in the dynamics of mean reversion over long time horizons, **it is therefore reasonable to conclude that a historically typical degree of mean reversion might be incorporated into volatility expectations.**

We conclude with a note of caution: the observation of mean reversion statistics requires a suitably long period of study, otherwise we might be calibrating on only a single market regime. The sizeable variation in the observed values of M and S over shorter measurement periods—and the greater stability at longer horizons—is provided for the interested reader in Appendix B.

VIX PREMIUM

“I hate to lose more than I love to win.”

- Jimmy Connors

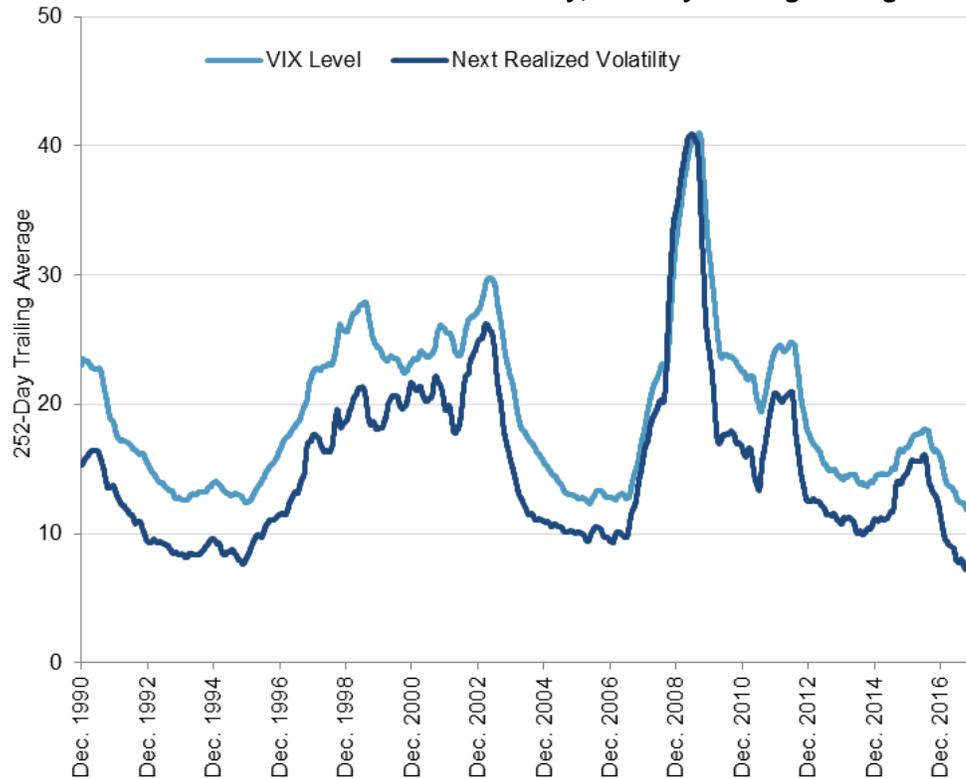
VIX has more often than not proved to be an over-estimate for future volatility, measured by next realized volatility.

Recall that our program to estimate where VIX is “expected to be” comprises two main steps. The previous section completed the first step; showing how a value for expected volatility over the next 30 days, *all else being equal*, can be derived from mean reversion. Henceforth for notational convenience, “**MR volatility**” will refer to the “anticipated volatility under mean reversion,” calculated by taking then-current recent volatility and adding a value equal to the (historically calibrated) speed of mean reversion, times the difference from the then-current recent volatility to its (historically calibrated) mean. **It remains to describe what differences are expected between MR volatility and the then-current VIX.**

Note that as well as making the resulting formulas simpler, breaking down the VIX estimation into a multi-stage process (first calculating MR volatility, and then the difference of that to the expected VIX level) allows us to gain a valuable insight into one of the more intriguing phenomena of options markets: the systemic premium in the volatility “implied” by options prices.

Certainly, VIX has more often than not proved to be an overestimate for future volatility. This was visible in Exhibit 2 for the first ten months of 2017, while Exhibits 10 and 11 demonstrate both the phenomenon and the advantage of using MR volatility in examining it over the long term. Exhibit 10 shows the historical degree of overestimation by comparing a 252-trading-day trailing average of VIX to the corresponding trailing average of next realized volatility. A comparison of these trailing averages reduces the dependence of the data on any one day or event. Exhibit 10 shows that a clear premium of around four to five points is visible in VIX most of the time.

Exhibit 10: VIX Versus Next Realized Volatility, 252-Day Trailing Average

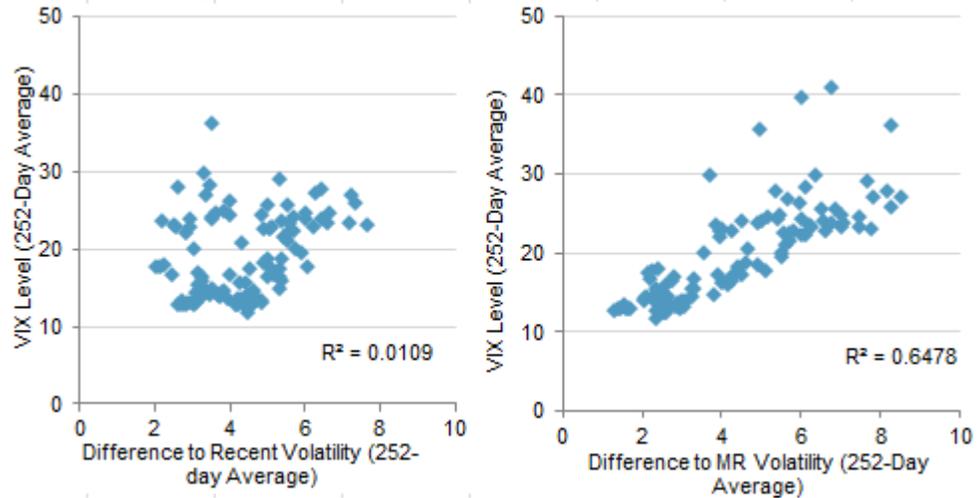


Source: S&P Dow Jones Indices LLC and CBOE. Data from Jan. 1, 1990, to Oct. 31, 2017. Chart is based on VIX levels and their corresponding S&P 500 recent volatility levels on each trading day. Past performance is no guarantee of future results. Chart is provided for illustrative purposes.

Exhibit 11 shows how the premium (or overestimate) in VIX appears to be more closely related to MR volatility than to the then-current realized volatility.⁷ Both charts in Exhibit 11 display a historical scatter plot of the level of VIX, versus the difference between VIX and a then-current measure: recent volatility in the case of the left-hand chart and MR volatility on the right. As with Exhibit 10, we used 252-day trailing averages in order to reduce the dependence on any particular day or event.

⁷ MR volatility has been calculated using a constant M=15, and S=27% for the full historical period.

Exhibit 11: VIX Overestimate Is More Closely Related to MR Volatility Than Recent Volatility



Source: S&P Dow Jones Indices LLC and CBOE. Data from Jan. 1, 1990, to Oct. 31, 2017. Chart is based on VIX levels and their corresponding S&P 500 recent volatility levels on each trading day. MR Volatility is calculated using $M = 15\%$ and $S = 27\%$. An average of the previous 252 values for VIX and the difference to recent volatility and MR volatility is taken on each day starting on Dec. 28, 1990. Dates are then selected at 60-day intervals in the period from December 1990 to October 2017. Past performance is no guarantee of future results. Chart is provided for illustrative purposes.

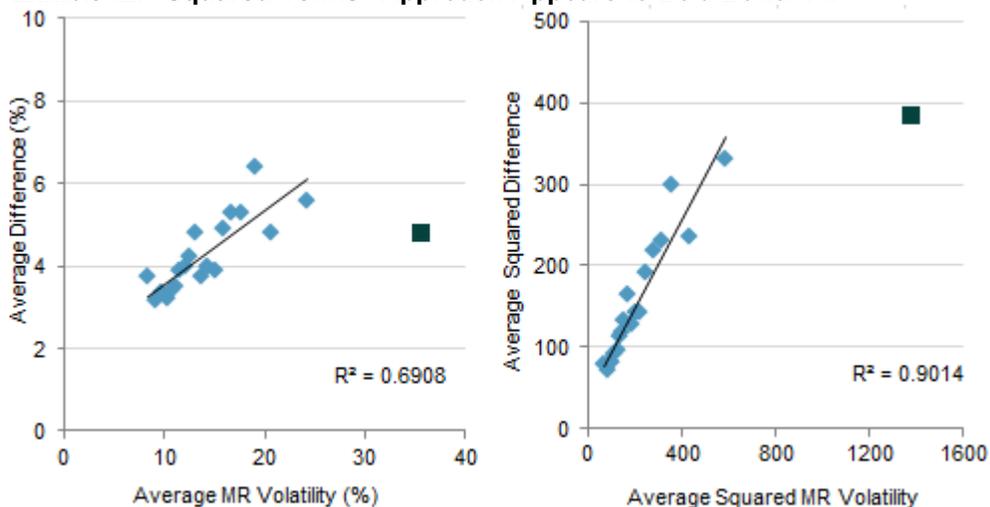
If we wish to examine what the typical “premium” in VIX depends on, then MR volatility might be a better place to start than recent volatility.

Exhibit 11 shows that if we wish to examine what the typical “premium” in VIX depends on, then MR volatility might be a better place to start than recent volatility. Indeed, the observed regression fit (R^2) of 0.64 in the right-hand chart suggests that a simple straight-line relationship may provide a reasonable estimate. We shall instead argue for the application of a slightly more complicated interaction, specifically between the *square* of VIX, and the *square* of MR volatility.⁸

Exhibit 12 provides evidence in support of such a “squared-terms” approach, using the same percentile-based approach used in Exhibits 4 and 7 to examine the average relationship between two sets of data. The left-hand chart compares the average MR volatility level to the average difference between VIX and MR volatility (“average difference”), bucketed by 5% ranges of MR volatility. The right-hand chart does the same thing with squared terms. That is to say, the right-hand chart compares the average of MR volatility *squared* to the average difference between VIX *squared* and MR volatility *squared*, within the same percentile bands as the left-hand chart. VIX squared minus MR volatility squared is referred to as the “squared difference.”

⁸ There are good, although complex, reasons for supposing the correct form of relationship between VIX and MR volatility is through their squared values. Simplifying somewhat, the level of VIX is constrained within certain arbitrage bands by the cost of a specific basket of options (and the expected cost of maintaining that basket over the life of the options). With a few assumptions, if appropriately maintained, that basket will deliver a payout equal to the future *variance* (volatility squared) of the asset underlying the options, minus the initial cost of establishing the basket. It is natural, therefore, to look for a premium in variance, opposed to volatility, terms. The interested reader is directed to Demeterfi et. al., “[More than you ever wanted to know about volatility swaps](#),” Goldman Sachs Quantitative Strategies Research Notes (March 1999), for a more in-depth exposition.

Exhibit 12: “Squared-Terms” Approach Appears to Be a Better Fit



Source: S&P Dow Jones Indices LLC and CBOE. Data from Jan. 1, 1990, to Oct. 31, 2017. Charts are based on VIX levels and their corresponding S&P 500 recent volatility levels and MR volatility levels on each trading day. MR Volatility is calculated using M = 15% and S = 27%. Past performance is no guarantee of future results. Charts are provided for illustrative purposes.

Historically, a better fit to the data is offered by assuming that any “premium” in VIX is determined by a linear relationship in *variance*, not volatility.

Note that in either chart, the 20th bucket (highlighted in dark blue and corresponding to days when MR volatility was in the highest 5% of values) appears as an outlier; it does not fit the wider pattern in either chart.

However, when excluding such outlying values, Exhibit 12 demonstrates that—at least historically and on average—a **better fit to the data is offered by assuming that any “premium” in VIX is determined by a linear relationship in *variance*, not volatility.** Equivalent analyses of the relative fit of squared versus non-squared variables for other VIX indices (for other equity markets, currencies, and so on) are provided in Appendix D, the majority of which support a similar conclusion.

For notational convenience, we therefore define a “**Variance Premium**” (**VP**) at each point in time as *the anticipated difference between the square of then-current MR volatility and the square of then-current VIX, based on a historical regression of the two squared variables.* We shall see that the VP provides us with the critical part of our third component in the VIX decomposition indicated earlier.

DECOMPOSING VIX, REVISITED

Combining the previous two sections allows us to provide an equation for where VIX is expected to be, assuming a continuation of historical norms and given a level of recent volatility—**this expectation is henceforth referred to as EVIX.** Specifically, EVIX is defined as follows:

$$EVIX = \sqrt{MR\ volatility^2 + VP}$$

Where, as before:

- **MR Volatility** = Recent Volatility + $S * (M - \text{Recent Volatility})$ is the “mean reversion volatility,” where the mean, **M**, and speed, **S**, are parameters of mean regression in recent volatility observed from historical data; and
- **VP** = $c * (\text{MR volatility})^2 + d$ is the “Variance Premium,” where **c** and **d** are constants observed by regressing the historical squared values of MR volatility to the difference between MR volatility squared and VIX squared.

In notational terms, we define the DTM on any day calculated as the then-current difference between VIX and EVIX.

We have now almost completed the program; the first three components of our decomposition sum to equal the “expected” VIX level, EVIX. The *difference* between EVIX and the actual VIX level provides the missing final component. In notational terms, we define the “Difference to Model” (DTM) on any day as the then-current difference between VIX and EVIX:

$$DTM = VIX - EVIX$$

Note that by construction, adding this difference to EVIX will give us VIX, and if we define a “Volatility Premium” and “Mean Reversion Adjustment” (MR Adjustment) according to the equations:

$$\text{Volatility Premium} = \sqrt{MR \text{ volatility}^2 + VP} - MR \text{ Volatility}$$

$$MR \text{ Adjustment} = MR \text{ Volatility} - \text{Recent Volatility}$$

Then, we have that on any day:

$$\mathbf{VIX = Recent Volatility + MR Adjustment + Volatility Premium + DTM.}$$

This gives an explicit form for our decomposition of VIX. In particular, note that with the exception of “recent volatility,” which itself is calculated from *recent* historical data, *the DTM is the only component not derived from long-term historical norms*; it therefore provides a direct candidate as a source of information regarding market sentiment or anticipated newsflow.

Note that the DTM reflects the degree to which VIX is “high” or “low” in comparison to a value for volatility that *already incorporates an expectation for mean reversion*. Thus, if we are to examine the predictive aspects of DTM, we should compare actual historical changes in volatility to the sum of the DTM and the MR adjustment. Since we shall examine this value in some detail, for notational purposes we define a “**VIX-Implied Change in Realized Volatility**” (VCR) as follows:

$$\mathbf{VCR = MR Adjustment + DTM}$$

A NUMERICAL EXAMPLE

An example of the possible use and calculation of the VCR may be given by returning to point “C” identified in Exhibit 3, corresponding to the levels of VIX and recent volatility as they might have been observed on July 18, 2016. Performing a historical regression on recent volatility in the S&P 500 from Jan. 27, 1928, to July 18, 2016, we find the mean reversion statistics $M = 15\%$ and $S = 30\%$. Applied to a then-current recent volatility of 19.7%, the MR adjustment was -1.4%, and so MR volatility was 18.3%.

Performing a historical regression on MR volatility squared and its difference to VIX squared on the available historical data from Jan. 2, 1990, to July 18, 2016, we find that $c = 0.60$ and $d = 26$ providing an EVIX level of 23.8. Since VIX stood at 12.4 on July 18, 2016, the DTM was -11.4. Combined with the MR adjustment of -1.4, the VCR was -12.8.

As it happens, over the subsequent 30 calendar days, there was an observed decrease in recent volatility, from 19.7% to 5.6%, a negative change equal to 14.1 percentage points—more than the 12.8 percentage point fall that VCR would have suggested, and more than the 1.4 point drop that that would have been expected under mean reversion alone.

THE PREDICTIVE RECORD OF THE VCR

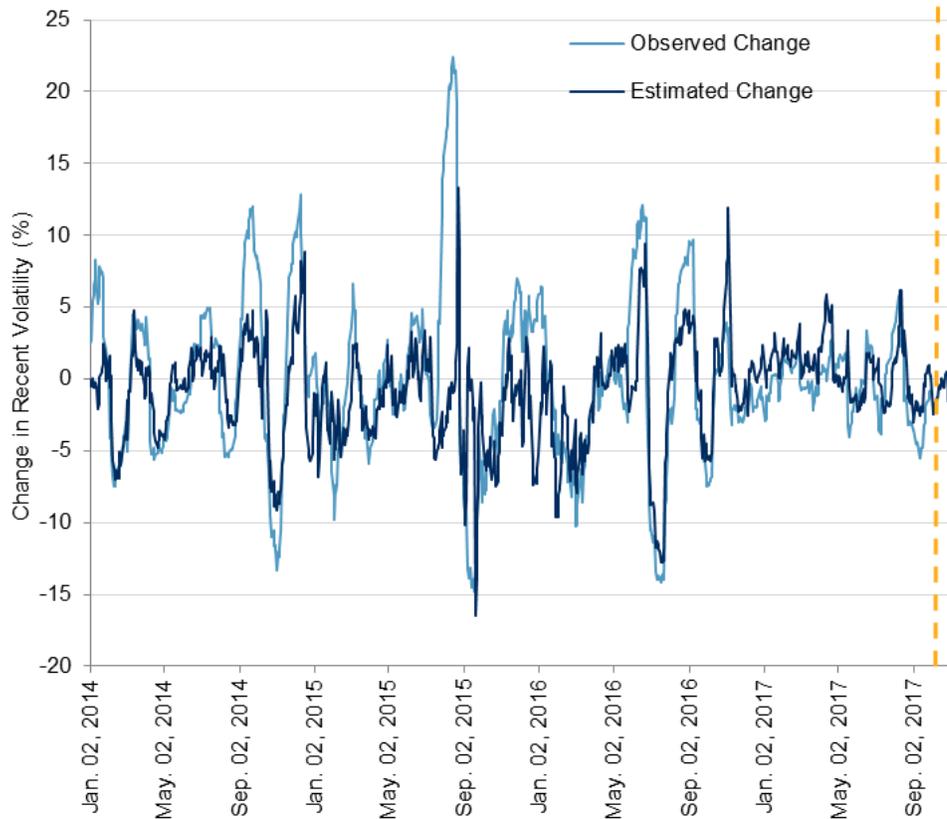
When examining the predictive power of our interpretation of VIX, one must exercise extreme caution; our analysis has relied on a historical dataset in order to explore putative relationships between VIX and realized volatility or derivations thereof. In testing the value of any resulting predictions, **we are testing a model on the very data that inspired the model and on which it was calibrated**. This is a version of so-called “look-ahead bias,” which can be mitigated, but not eliminated. The authors were cognizant of the historical patterns and norms in volatility even when embarking on this study; a different history might have resulted in an entirely different model, as well as different calibrations for that model.

Duly cautioned, Exhibit 13 shows the historically observed changes in recent volatility compared to those “anticipated” by the VCR at the start of the period. For example: the value for the “observed change” series corresponding to Jan. 2, 2014, is the change in realized volatility from Jan. 2, 2014, to Feb. 2, 2014, a value that was only known on the later date. The value for the VCR series would have been calculable on the earlier date. In order to diminish the degree of “look-ahead bias,” the values for the VCR series were calculated using only the historical data that would have been available *at that time*.⁹

Our analysis has relied on a historical dataset in order to explore putative relationships between VIX and realized volatility or derivations thereof.

⁹ The full detail of how this was achieved may be found in Appendix C.

Exhibit 13: Assessing the Predictive Power of VCR



Source: S&P Dow Jones Indices LLC and CBOE. Data from Jan. 2, 2014, to Oct. 30, 2017. Chart is based on the VCR and the differences between recent volatility and next realized volatility in the S&P 500, as calculated on each trading day when possible. MR Volatility and the VP are calculated in such a way that look-ahead bias is removed. Past performance is no guarantee of future results. Chart is provided for illustrative purposes.

The levels of, and the movements in, the two series of Exhibit 13 are encouragingly similar, which essentially shows that the VCR would have functioned as an **always imperfect but nonetheless meaningful** prediction for observed future changes in volatility. We emphasize the importance of the fact that the dark blue “VCR” series extends beyond the latest date of the light blue “observed change” series; the level of the former series is available 30 days before the latter.

VCR would have functioned as an always imperfect, but otherwise meaningful prediction for the observed future changes in volatility.

For a longer-term perspective, Exhibit 14 compares the average observed 30-day change in recent volatility to the VCR at the start of the period, for all trading days between Dec. 21, 1999, and Sept. 29, 2017. At each point, the VCR was calibrated using only historical data available at that point in time. Exhibit 14 shows the average changes in recent volatility based on various ranges for the starting level of recent volatility and the VCR. For example, when recent volatility was between 25% and 30% at the start of the period and VCR was less than -2, then over the subsequent 30 calendar days, the average change in recent volatility was a decrease of 4.40%. Note that some of the entries in the table reflect relatively “rare” circumstances; asterisks denote when the sample contained fewer than 21

historical observations. Note also that we limited our observations to those when realized volatility was no greater than 30%, in line with our remarks after Exhibit 12 observing that the VP followed non-standard dynamics when recent volatility was in the highest 5% of values (corresponding to recent volatility at any level higher than 30.8%).

Exhibit 14: Average Change in Recent Volatility for Given Levels of Recent Volatility and VCR

RECENT VOLATILITY (%)	VCR					
	< -2	-2 to -1	-1 to 0	0 to 1	1 to 2	> 2
0 to 9	-0.95*	0.59	0.61	1.23	1.98	5.05
9 to 10	-0.45	1.02	2.84	2.79	2.62	3.78
10 to 12	-0.63	0.30	1.55	0.94	1.85	4.37
12 to 14	-2.78	-0.05	-0.18	2.97	1.76	4.89
14 to 17	-3.42	0.00	0.74	1.28	2.65	4.38
17 to 20	-4.85	0.56	1.58	1.00	0.89	3.84
20 to 25	-1.17	0.00	0.40	1.56	2.73	1.27
25 to 30	-4.40	3.55*	3.89*	5.24*	4.19*	3.58

Source: S&P Dow Jones Indices LCC and CBOE. Data from Dec. 21, 1999 to Sept. 29, 2017. Table is based on the change in recent volatility over 30 calendar days. The data is separated into categories based on the value of recent volatility at the start of the each 30-calendar-day period and depending on the value of VCR. MR Volatility and the VP are calculated in such a way that look-ahead bias is removed. *The sample comprised fewer than 21 historical observations. Past performance is no guarantee of future results. Table is provided for illustrative purposes.

Exhibit 14 shows that VCR has been a reasonable indicator of future changes in recent volatility, particularly when it differs significantly from zero. It is important to emphasize that although the VCR is a far from perfect predictor, it is considerably better than several alternatives. Exhibit 15 compares the accuracy of the VCR in anticipating the 30-day change in realized volatility in comparison to three simpler alternatives: the first using recent volatility as an estimate for future volatility, the second using the MR volatility (calibrated on the historical data that was known at the time), and the final using the then-current VIX as an outright prediction.

In order to construct Exhibit 15, over each date between Dec. 21, 1999, and Sept. 30, 2017, but ignoring days when recent volatility exceeded 30%,¹⁰ we compared the *actual* change in recent volatility to that “predicted” by each of the four alternatives, and measured their absolute error. Exhibit 15 shows the median, average, as well as 25th and 75th percentiles of these errors.

¹⁰ See remarks following Exhibit 12 on the exclusion of outliers.

Exhibit 15: Average Error of the VCR in Comparison to Naïve Alternatives

ABSOLUTE DIFFERENCE	VCR	RECENT VOLATILITY	MR VOLATILITY	VIX
Median	2.31	3.06	3.02	4.62
Average	3.58	4.25	4.08	5.27
25 th Percentile	1.09	1.37	1.47	2.65
75 th Percentile	4.45	5.66	5.20	7.01

Source: S&P Dow Jones Indices LLC and CBOE. Data from Dec. 21, 1990, to Sept. 30, 2017. Past performance is no guarantee of future results. Table is provided for illustrative purposes.

As Exhibit 15 shows, the average and median absolute error of the VCR in predicting future changes in recent volatility is lower than the other models, but none are particularly accurate predictions. We suspect that although improvements to the accuracy of the VCR might be possible, there is a natural limit to the degree of accuracy of any such predictions; not only because of the ever-present risk of surprise events, but also because—as the next example demonstrates—even when a well-telegraphed risk manifests, the impact may be different than anticipated.

While the VCR is a far from perfect prediction for subsequent changes in recent volatility, it is better than naïve alternatives

HOW VIX GETS PREDICTIONS WRONG – AN INSTRUCTIVE EXAMPLE

“Never make predictions, especially about the future.”
- Casey Stengel

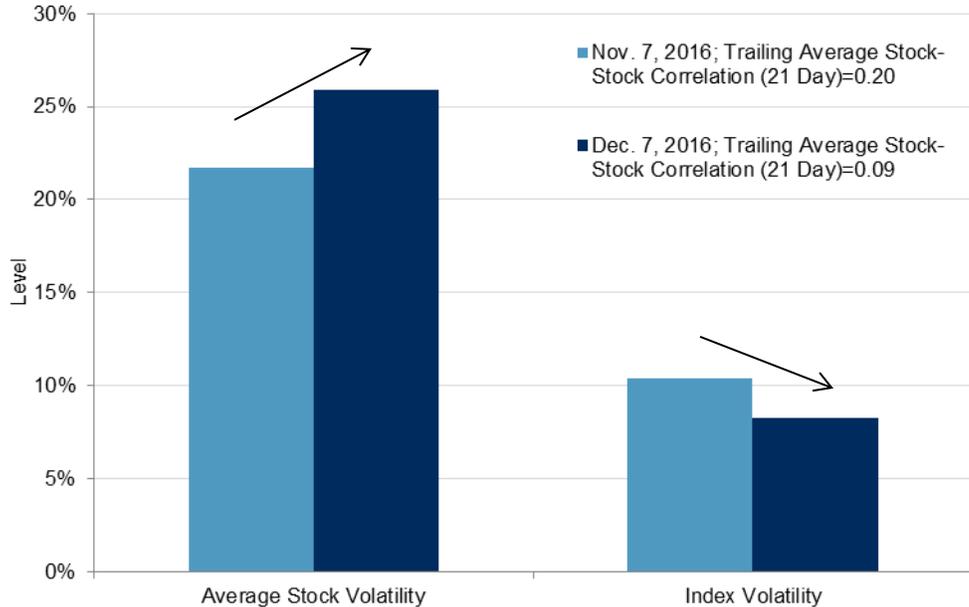
At the close of the trading day on Nov. 7, 2016, VIX stood at 18.7, recent volatility was 10.5%, and the VCR equaled 4.3. This would suggest that the market was anticipating a potentially significant increase in volatility. In fact, the high VIX was likely reflective of the uncertainty felt by market participants about the next day’s U.S. presidential election. Despite the results of that election defying consensus expectations, volatility actually *decreased* from its level just prior to the election result, recording a level of 8.4% when measured 30 days later.

If the particular circumstance of the U.S. election provides an example of where the predictive content of VIX proved misleading, the event also offers **a valuable lesson as to how VIX can “get things wrong.”** Recall that the S&P 500 is itself a *portfolio* of stocks. The volatility of the S&P 500, therefore, is dependent on two factors: the volatility of the constituent stocks *and* their correlations.¹¹ Hence, if correlations fall, it is possible for individual stock volatilities to rise even as portfolio volatility falls. Indeed, a fall in correlations may explain the apparent disconnect between expected and observed volatility levels around the 2016 U.S. presidential election. Exhibit 16 shows the average 21-day trailing annualized S&P 500 *stock*

¹¹ For more details on the interaction between volatility and correlation, please see Edwards, Tim and Craig J. Lazzara, “[The Landscape of Risk](#),” (December 2014).

volatility as well as the volatility of the S&P 500, as measured on Nov. 7 2016, and Dec. 7, 2016—one month later.

Exhibit 16: Changes to Recent Volatility of the S&P 500 Versus Changes to the Average S&P 500 Stock Volatility Around the 2016 U.S. Presidential Election



The fact that a decrease in correlations meant that index volatility fell, even though individual stock risk rose, provides an example of how VIX may essentially get risk right, but correlations wrong.

Source: S&P Dow Jones Indices LLC. Data from Nov. 7, 2016, to Dec. 7, 2016. Chart is based on the volatility of the S&P 500, and the average stock volatility in the S&P 500, as calculated over the prior 21 trading days. Past performance is no guarantee of future results. Chart is provided for illustrative purposes. See “[At the Intersection of Diversification, Volatility and Correlation](#)”; Edwards and Lazzara; S&P DJI Research (2014), for an explanation of how average stock-stock correlation is calculated.

Exhibit 16 shows that average stock volatility increased after the election by several percentage points. However, since correlations fell considerably, index volatility actually declined. A potential explanation for this result is that market participants took the view that different companies and sectors would be affected differently by the election result. Either way, the fact that a decrease in correlations meant that index volatility fell, even though individual stock risk rose, **provides an example of how VIX may essentially get risk right, but correlations wrong.**

APPLICATION TO VIX INDICES AROUND THE WORLD

While we leave the details to the appendix, **similar results may be found for a range of indices around the globe that use the VIX methodology.** In particular, applying the methods used in previous sections to various equity, currency, commodity, and fixed income markets provides a geographical and asset-class-based (if not temporal) “out-of-sample” test for the theory developed in this paper. The application of our techniques to these markets shows that (in summary):

- 1) Volatility in multiple markets has shown a tendency **for mean reversion that satisfies, on average, the form:**

Next Realized Volatility

$$\approx \text{Recent Volatility} + S * (M - \text{Recent Volatility})$$

Given the bank of literature highlighting the tendency of volatility to revert to its mean across various markets, this result is neither surprising nor novel. The main exception to the rule is the CBOE VIX of VIX (VVIX) volatility measure, which is a volatility-of-volatility measure that uses VIX options to derive an implied volatility for VIX. The evidence for mean regression of the form provided above is by far the weakest in the case of the VVIX.

- 2) Additionally, **our results regarding the form of a premium that is linear in variance terms (volatility squared) appear to hold true across many markets.**¹² In most of the markets we examined, there is a stronger linear relationship between MR volatility squared and its difference to VIX squared, than in the relationship in non-squared terms. Only VVIX shows a stronger relationship between unsquared terms. Interestingly, the VVIX is also the only index for which the slope of the regression line is *negative*; the 0.99 R-squared figure for the VP in VVIX is the highest of all the indices we examine, and pertains to a negative slope of regression. This may suggest that **whatever premium might be available to the sellers of options in general markets, sellers of VIX options might not capture.**
- 3) As with VIX, the predictive power and reliability of the regression statistics in determining where VIX is “expected” to be **breaks down at particularly high levels of realized volatility** (in particular, the highest 5% of readings for realized volatility appear to demonstrate different dynamics).
- 4) Finally, the VCR is a useful prediction for future changes in recent volatility in many markets. Indeed, a close relationship, as in Exhibit 13, is observed across many indices, therefore suggesting that the theory developed in this paper has the potential to offer useful insights around the world.

¹² For more discussion on the VP, please see Bekaert and Hoerova, “[The VIX, the Variance Premium and Stock Market Volatility](#)” Journal of Econometrics (December 2014).

CONCLUSION

Without sophisticated mathematics or intricate models of market behavior, there are fairly simple ways to decode the information embodied in a given VIX level. Principally, market participants would have been well served, at least from a historical perspective, to account for recent volatility when reading VIX levels and comparing the current level to an indication of where VIX *ought to be* given its historic norms.

Several components explain the average relationship between VIX and the realized volatility of the S&P 500. These are, explicitly, the *recent* level of realized volatility, plus an anticipated *mean reversion* in volatility, plus a *premium* in VIX, itself seemingly *scaling according to anticipated variance* (not anticipated volatility). Using these components to estimate VIX based on historical patterns and comparing this estimate to observed VIX levels allows us to glean indications about market expectations for the future path of volatility, or how the supply and demand for options on the underlying S&P 500 is evolving.

In particular, our analysis showed how a so-called “VIX-Implied Change in Realized Volatility” (VCR) can be calculated, and demonstrated that this VCR performs reasonably well at approximating monthly movements in realized volatility across a range of VIX indices. Nonetheless, simply being able to understand what is “implied” in “implied volatility” should not be confused with being able to make perfect predictions. This report provides a guide to reading VIX; it remains for the reader to choose what to do with the information it provides.

BIBLIOGRAPHY

“The CBOE Volatility Index –VIX®”, White Paper, CBOE.

“At the Intersection of Diversification, Volatility and Correlation,” Edwards and Lazzara, 2014.

“The Landscape of Risk,” Edwards and Lazzara, 2015.

“The VIX, the Variance Premium and Stock Market Volatility,” Bekaert and Hoerova, *Journal of Econometrics*, December 2014.

“What good is a volatility model,” Engle and Patton, *Quantitative finance* 1, no. 2 (2001): 237-245.

“More than you ever wanted to know about volatility swaps,” Demeterfi et. al., Goldman Sachs Quantitative Strategies Research Notes, March 1999.

APPENDIX A: COMPUTING RECENT VOLATILITY

Recent volatility in the S&P 500 is calculated on any trading day according to the formula:

$$\text{Recent volatility} = \sqrt{\frac{252}{N} \times \sum_{t=1}^{N-1} \ln\left(\frac{p_{t+1}}{p_t}\right)^2}$$

Where p_1, p_2, \dots, p_N are the daily closing levels of the S&P 500 price index on consecutive trading days, the final closing level p_N corresponding to the current date and the first closing level p_1 corresponding to the last trading day that was more than 30 calendar days prior to the current date.

APPENDIX B: THE VARIATION IN MEAN REVERSION STATISTICS FOR S&P 500 RECENT CALCULATED OVER DIFFERENT TIME HORIZONS.

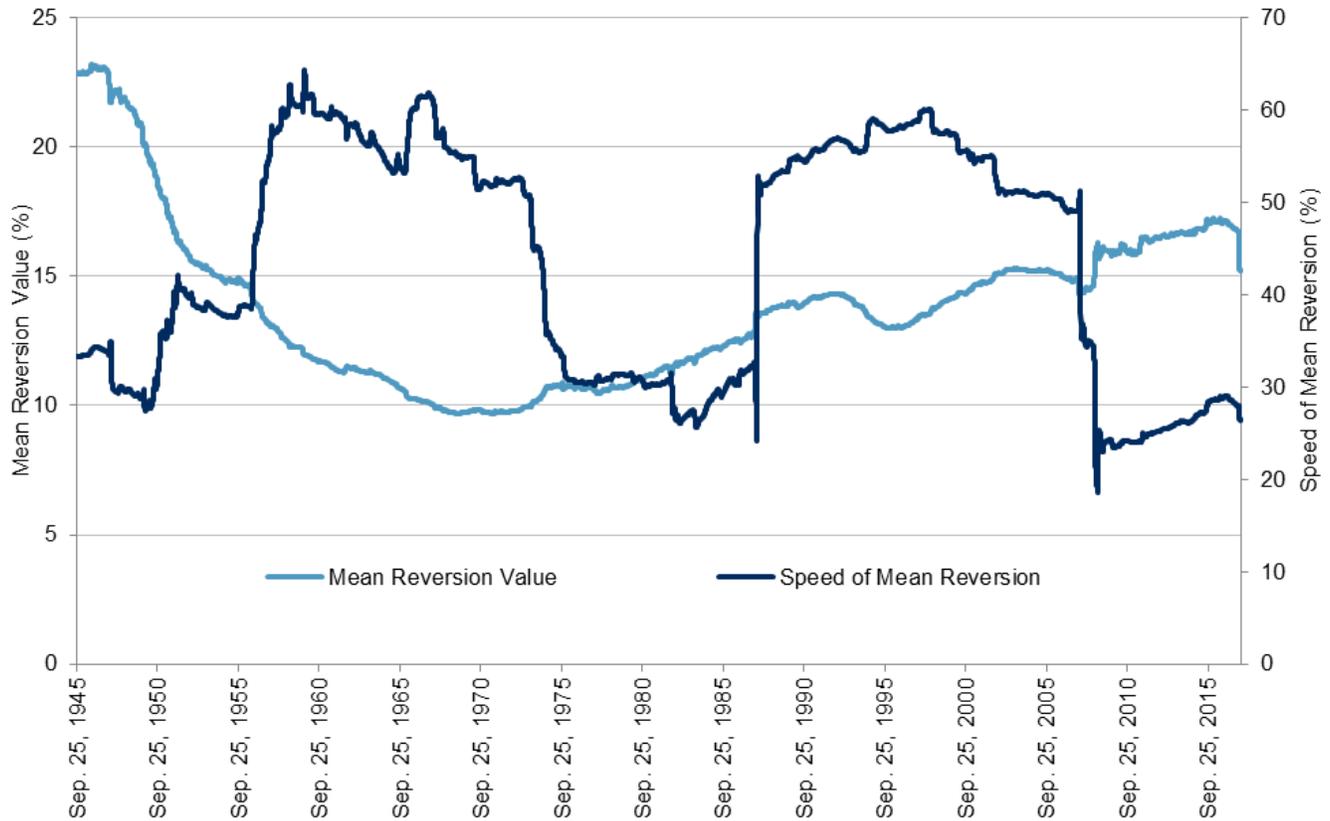
Exhibit 17 shows the average, minimum, maximum, and median values of M and S when calculated on a trailing basis on each day in the period from January 1928 to September 2017, over various different time horizons. Specifically, over rolling 1-, 3-, 5-, 10-, 20-, 30-, and 50-year horizons, we regressed next realized volatility on recent volatility using daily S&P 500 closing price data. M and S are calculated for each regression, based on the equation provided in the section titled “Mean Reversion in Volatility.”

TIME HORIZON (YEARS)	SPEED (%)				MEAN (%)			
	AVERAGE	MINIMUM	MAXIMUM	MEDIAN	AVERAGE	MINIMUM	MAXIMUM	MEDIAN
1	79.49	-165.85	214.14	77.77	16.48	-532.53	21813.67	13.12
3	58.33	-30.45	121.35	58.34	15.57	-733.46	931.27	13.21
5	49.53	-25.49	94.08	47.99	15.23	-1652.83	1837.62	13.75
10	44.96	13.73	84.86	42.61	14.73	8.79	29.70	13.60
20	43.57	18.52	64.38	49.04	13.86	9.68	23.21	13.79
30	42.02	24.86	55.66	42.83	13.40	10.56	19.26	12.85
50	39.33	23.28	52.28	40.87	13.16	11.71	15.82	12.91

Source: S&P Dow Jones Indices LLC. Data from Jan. 3, 1928, to Sept. 29, 2017. Past performance is no guarantee of future results. Table is provided for illustrative purposes.

APPENDIX B.1: MEAN REVERSION DYNAMICS IN THE S&P 500 OVER ROLLING 20-YEAR PERIODS

Exhibit 18: Mean Reversion Dynamics in the S&P 500



Source: S&P Dow Jones Indices LLC. Data from Sept. 25, 1945, and Sept. 29, 2017. Chart shows the values of M and S after regressing next realized volatility on recent volatility over rolling 20-year windows. S&P 500 closing prices between Jan. 3, 1928, and Oct. 29, 2017, are used to calculate recent volatility and next realized volatility. Past performance is no guarantee of future results. Chart is provided for illustrative purposes.

APPENDIX C: REMOVING THE LOOK-AHEAD BIAS IN EVIX.

On each date between Jan. 27, 1928, and Sept. 29, 2017, recent volatility and next realized volatility figures were calculated based on closing S&P 500 price return changes over 30-day periods. We then ran a regression of next realized volatility on recent volatility between Jan. 27, 1928, and the date in question to obtain estimates for the mean reversion dynamic. The first date on which the mean reversion variables M and S were estimated was Jan. 2, 1990, which is the date on which the first closing value of VIX was reported. MR volatility squared and the VP were computed on each day from Jan. 2, 1990, to Sept. 29, 2017. In order to avoid estimating the relationship between VP and MR volatility squared using days with elevated levels of volatility, on days where recent volatility exceeded 30%, we replaced MR volatility squared values and VP figures with their respective averages up to the date in question.

The final step is to regress VP on MR volatility squared between Jan. 2, 1990, and the date in question. The first date on which this regression is run is Dec. 21, 1999, such that the fewest number of days used in the regressions is 2,520 days, or 10 years. The VCR for each date was then calculated according to formula provided in the “Decomposing VIX, Revisited” section. For October 2017, we used the estimated parameters from Sept. 29, 2017, to transform recent volatility into EVIX on each date after Sept. 30, 2017, up to Oct. 31, 2017. The values of VCR in October 2017 therefore correspond to a predicted path of recent volatility, assuming a continuation of the observed relationships up to Sept. 29, 2017.

APPENDIX D: GLOBAL ESTIMATES FOR VIX

As a reminder, we assume that EVIX takes the following form:

$$EVIX = \sqrt{MR \text{ volatility}^2 + VP}$$

Where, as before:

- **MR Volatility** = Recent Volatility + $S * (M - \text{Recent Volatility})$ is the “mean reversion volatility,” where the mean, **M**, and speed, **S**, are parameters of mean regression in recent volatility observed from historical data.
- **VP** = $c * (MR \text{ volatility})^2 + d$ is the “Variance Premium,” where **c** and **d** are constants observed by regressing the historical squared values of MR volatility to the difference between MR volatility squared and VIX squared, based on local averages within five percentile bands and excluding all data where MR volatility was in the highest 5% of observations.

Exhibit 19 outlines the estimated parameters (M, S, c, and d) used to compute EVIX for a variety of indices based on VIX methodology. All values are computed across each index’s full history, with M and S being derived from the closing price levels of the underlying index for each VIX index; for example, we use VIX as the underlying index of CBOE VIX of VIX. The R² figures are for the locally averaged statistics – taken within all 5% bands for mean reversion and excluding the highest 5% in the case of the VP

Exhibit 19: Global Estimates for VIX

INDEX	UNDERLYING INDEX FOR MR AND VP	M (%)	S (%)	MR R ²	c	d	VP R ²
VIX	S&P 500	15	27	0.99	0.5	39	0.90
S&P/ASX 200 VIX	S&P/ASX 200	15	24	0.97	0.79	-31	0.91
VSTOXX®	Euro STOXX 50	21	26	0.99	0.27	64	0.66
S&P/TSX 60 VIX	S&P/TSX 60	14	27	0.99	0.23	60	0.55
HSI Volatility Index	Hong Kong Hang Seng Index	20	18	0.99	0.4	-2	0.94
S&P/JPX JGB VIX®	S&P 10-Year JGB Futures Index (TR)	2	27	0.98	0.62	2	0.88
CBOE/CME FX Euro Volatility	EUR/USD Spot Rate	10	26	0.97	0.37	-5	0.77
CBOE/CME FX GBP Volatility	GBP/USD Spot Rate	10	28	0.97	0.54	-19	0.84
CBOE/CME FX Yen Volatility	Yen/USD Spot Rate	10	51	0.91	0.64	-22	0.89
CBOE Gold ETF Volatility	LBMA Gold Price PM	17	33	0.97	0.44	-3	0.82
CBOE VIX of VIX	VIX	109	85	0.33	-0.05	-3651	0.04

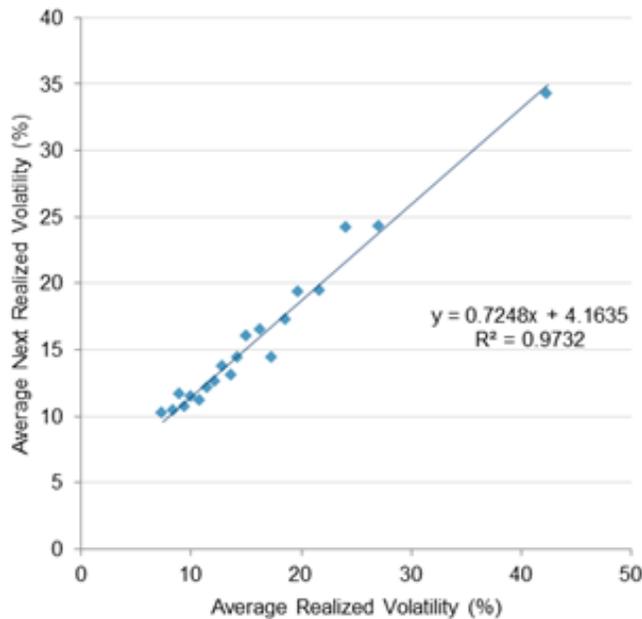
Source: S&P Dow Jones Indices LLC. Data as of October 2017. Past performance is no guarantee of future results. Table is provided for illustrative purposes.

S&P/ASX 200 VIX

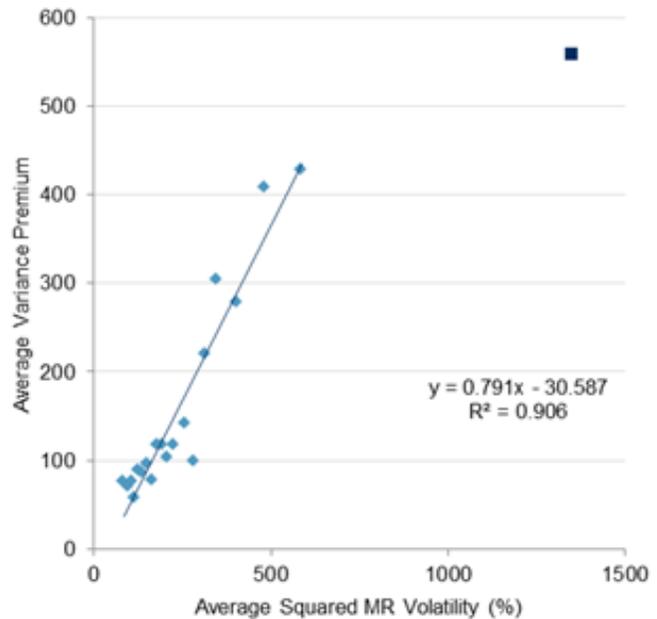
Recent volatility in the S&P/ASX 200 has shown a strong tendency to mean revert; the goodness-of-fit (R-squared) is 0.97 for the positive linear relationship between average recent volatility and next realized volatility. A similarly strong R-squared value of 0.91 is observed in the linear relationship between the squared MR volatility and the VP; this compares to an R-squared figure of 0.71 for the

linear relationship between MR volatility and the difference (between the S&P/ASX 200 VIX and MR volatility).

Exhibit 20: MR and VP of the S&P/ASX 200 VIX

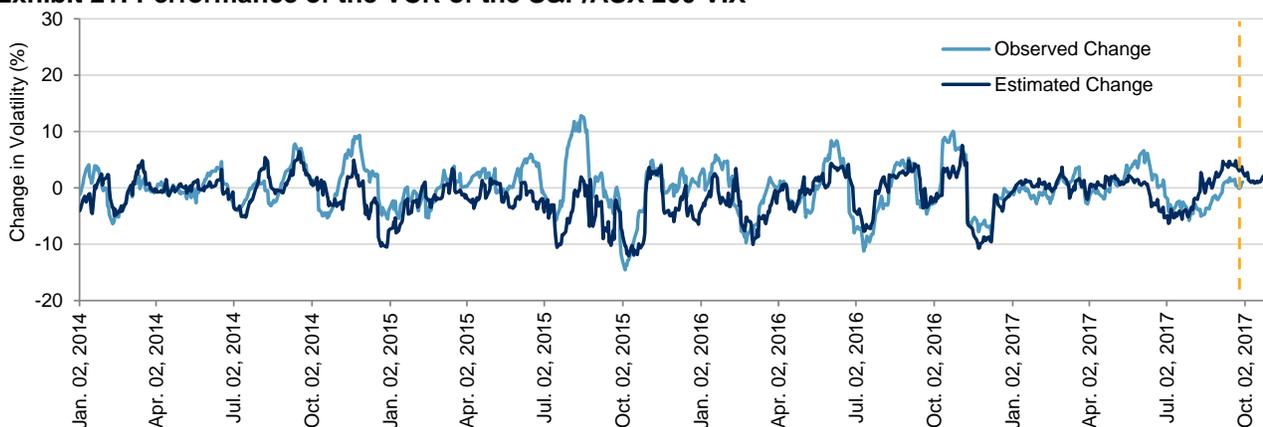


Source: S&P Dow Jones Indices LLC. Data from Feb. 1, 2008, to Sept. 29, 2017, using closing price levels in the index from Jan. 2, 2008, to Oct. 31, 2017. Recent volatility and next realized volatility for the S&P/ASX 200 were calculated on each day. Recent volatility and next realized volatility values were then evenly grouped into 20 buckets depending on the percentile rank of recent volatility across the whole period. The chart shows the average recent volatility and the average next realized volatility in each group. Past performance is no guarantee of future results. Chart is provided for illustrative purposes.



Source: S&P Dow Jones Indices LLC and CBOE. Data from Feb. 1, 2008, to Sept. 29, 2017. The MR volatility values suffer from look ahead bias. MR volatility is calculated using a constant $M=15\%$, $S=24\%$ for the period. MR volatility was squared and we calculated the VP on each date. These values were then evenly grouped into 20 buckets depending on the percentile rank of recent volatility across the whole period. The chart shows the average MR volatility and the average VP in each group. The line of best fit was calculated using only the first 19 buckets; it ignored days when recent volatility is in the highest 5% of all readings. Past performance is no guarantee of future results. Chart is provided for illustrative purposes.

Exhibit 21: Performance of the VCR of the S&P/ASX 200 VIX

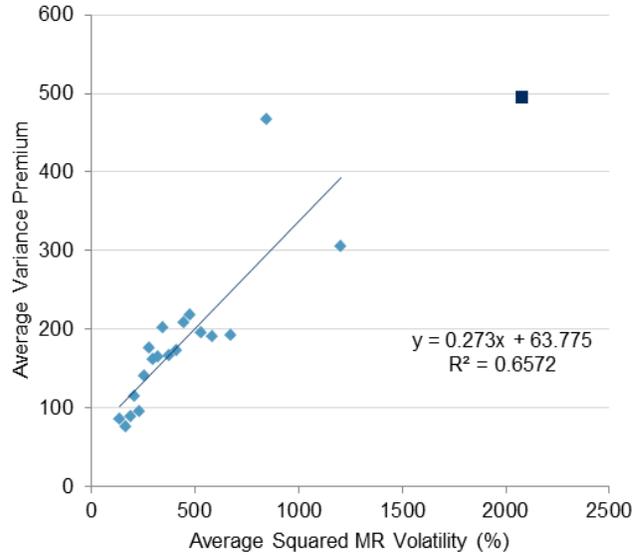
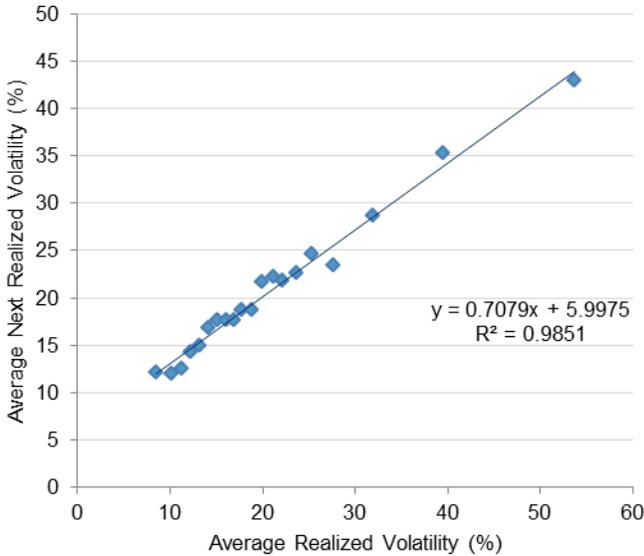


Source: S&P Dow Jones Indices LLC and CBOE. Data from Jan. 2, 2014, and Sept. 29, 2017. VCR is calculated on each date and suffers from look-ahead bias prior to Sep. 29, 2017. Values for October 2017 are calculated assuming historical relationships are unchanged. These values for VCR compared to changes in recent volatility in the S&P/ASX 200 over the corresponding 30-day period. Past performance is no guarantee of future results. Chart is provided for illustrative purposes.

VSTOXX

Recent volatility in the Euro STOXX 50 has shown a strong tendency to mean revert; the goodness-of-fit (R-squared) is 0.99 for the linear relationship between average recent volatility and next realized volatility. A reasonably strong R-squared value of 0.66 is observed in the linear relationship between the squared MR volatility and the VP; this compares to an R-squared figure of 0.23 for the linear relationship between MR volatility and the difference (between the VSTOXX and MR volatility).

Exhibit 22: MR and VP of the VSTOXX



Source: Eurex and S&P Dow Jones Indices LLC. Data from Feb. 2, 1999, to Sept. 29, 2017. Recent volatility and next realized volatility for the Euro STOXX 50 were calculated on each day, using closing price levels in the index from Jan. 4, 1999, to Oct. 31, 2017. Recent volatility and next realized volatility values were then evenly grouped into 20 buckets depending on the percentile rank of recent volatility across the whole period. The chart shows the average recent volatility and the average next realized volatility in each group. Past performance is no guarantee of future results. Chart is provided for illustrative purposes.

Source: Eurex and S&P Dow Jones Indices LLC. Data from Feb. 2, 1999, to Sept. 29, 2017. MR volatility is calculated using a constant $M=21%$, $S=26%$ and suffers from look-ahead bias. MR volatility was squared and we calculated the VP on each date. These values were then evenly grouped into 20 buckets depending on the percentile rank of recent volatility across the whole period. The chart shows the average MR volatility and the average VP in each group. The line of best fit was calculated using only the first 19 buckets; it ignored days when recent volatility is in the highest 5% of all readings. Past performance is no guarantee of future results. Chart is provided for illustrative purposes.

Exhibit 23: Performance of the VCR of the VSTOXX

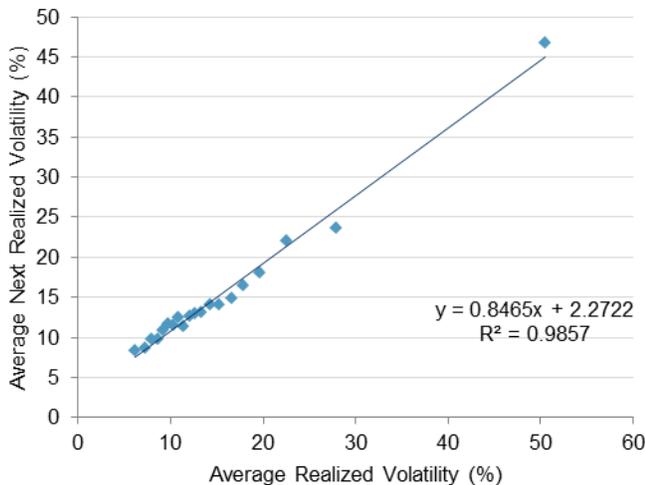


Source: Eurex. Data from Jan. 2, 2014, to Sept. 29, 2017. VCR is calculated on each date and suffers from look-ahead bias prior to Sept. 29, 2017. Values for October 2017 were calculated assuming historical relationships were unchanged. These values for VCR compared to changes in recent volatility in the Euro STOXX 50 over the corresponding 30-day period. Past performance is no guarantee of future results. Chart is provided for illustrative purposes.

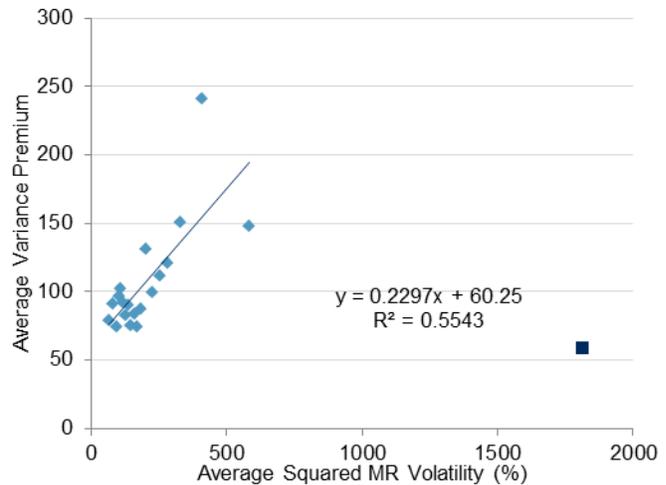
S&P/TSX 60 VIX

Recent volatility in the S&P/TSX 60 has shown a strong tendency to mean revert; the goodness-of-fit (R-squared) is 0.99 for the linear relationship between average recent volatility and next realized volatility. A reasonably strong R-squared value of 0.59 is observed in the linear relationship between the squared MR volatility and the VP; this compares to an R-squared figure of 0.01 for the linear relationship between MR volatility and the difference (between the S&P/TSX 60 VIX and MR volatility).

Exhibit 24: MR and VP of the S&P/TSX 60 VIX



Source: S&P Dow Jones Indices LLC. Data from Jan. 4, 2005, to Sept. 29, 2017. Recent volatility and next realized volatility for the S&P/TSX 60 were calculated on each day in the period using closing price levels in the index from Dec. 1, 2004, to Oct. 31, 2017. Recent volatility and next realized volatility values were then evenly grouped into 20 buckets depending on the percentile rank of recent volatility across the whole period. The chart shows the average recent volatility and the average next realized volatility in each group. Past performance is no guarantee of future results. Chart is provided for illustrative purposes.



Source: S&P Dow Jones Indices LLC/CBOE. Data from Jan. 4, 2005, to Sept. 29, 2017. MR volatility was calculated using a constant M=14%, S=27% for the period and suffers from look-ahead bias. MR volatility was squared and we calculated the VP on each date. These values were then evenly grouped into 20 buckets depending on the percentile rank of recent volatility across the whole period. The chart shows the average MR volatility and the average VP in each group. The line of best fit was calculated using only the first 19 buckets; it ignored days when recent volatility is in the highest 5% of all readings. Past performance is no guarantee of future results. Chart is provided for illustrative purposes.

Exhibit 25: Performance of the VCR of the S&P/TSX 60 VIX



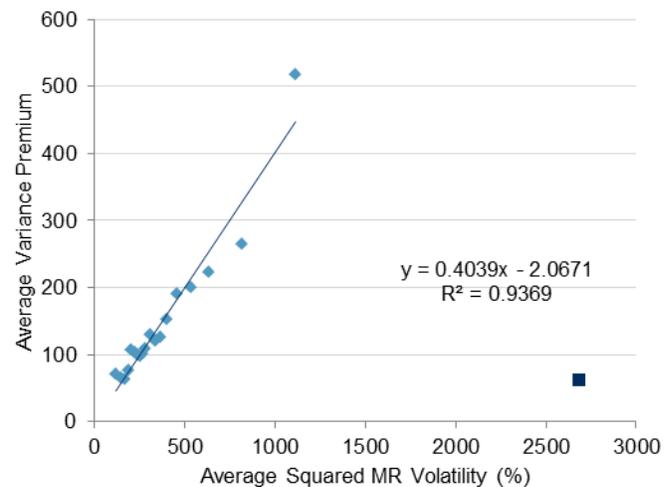
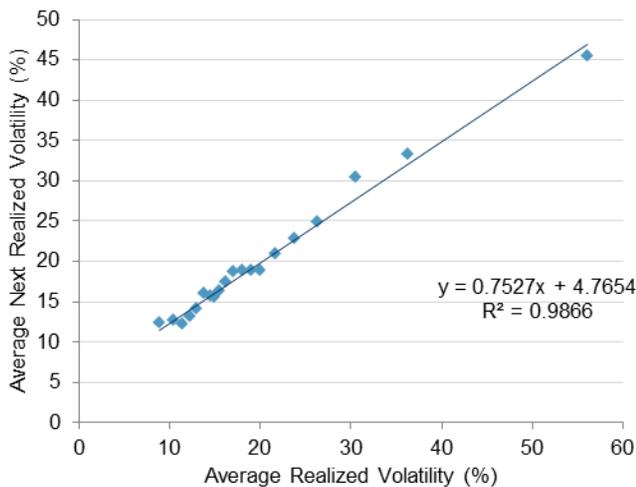
Source: S&P Dow Jones Indices LLC and CBOE. Data from Jan. 2, 2014, to Sept. 29, 2017. VCR is calculated on each date and suffers from look-ahead bias prior to Sept. 29, 2017. Values for October 2017 were calculated assuming historical relationships were unchanged. These

values for VCR compared to changes in recent volatility in the S&P/TSX 60 over the corresponding 30-day period. Past performance is no guarantee of future results. Chart is provided for illustrative purposes.

HSI Volatility Index

Recent volatility in the Hang Seng Index has shown a strong tendency to mean revert; the goodness-of-fit (R-squared) is 0.99 for the linear relationship between average recent volatility and next realized volatility. A similarly strong R-squared value of 0.93 is observed in the linear relationship between the squared MR volatility and the VP; this compares to an R-squared figure of 0.78 for the linear relationship between MR volatility and the difference (between the Hang Seng Volatility Index and MR volatility).

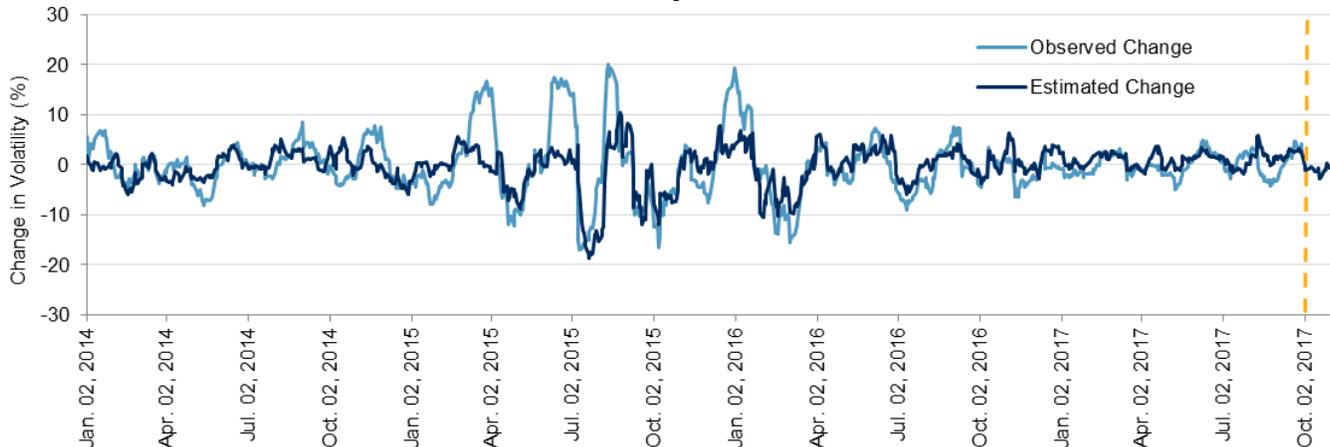
Exhibit 26: MR and VP of the HSI Volatility Index



Sources: Hang Seng and S&P Dow Jones Indices LLC. Data from Jan. 2, 2001, to Sept. 29, 2017. Recent volatility and next realized volatility for the Hong Kong Hang Seng Index were calculated on each day using closing price levels in the index from Nov. 29, 2000, to Oct. 31, 2017. Recent volatility and next realized volatility values were then evenly grouped into 20 buckets depending on the percentile rank of recent volatility across the whole period. The chart shows the average recent volatility and the average next realized volatility in each group. Past performance is no guarantee of future results. Chart is provided for illustrative purposes.

Source: S&P Dow Jones Indices LLC and CBOE. Data from Jan. 2, 2001, to Sept. 29, 2017. MR volatility is calculated using a constant M=20%, S=18% for the period and suffers from look ahead bias. MR volatility was squared and we calculated the VP on each date. These values were then evenly grouped into 20 buckets depending on the percentile rank of recent volatility across the whole period. The chart shows the average MR volatility and the average VP in each group. The line of best fit was calculated using only the first 19 buckets; it ignored days when recent volatility is in the highest 5% of all readings. Past performance is no guarantee of future results. Chart is provided for illustrative purposes.

Exhibit 27: Performance of the VCR of the HSI Volatility Index



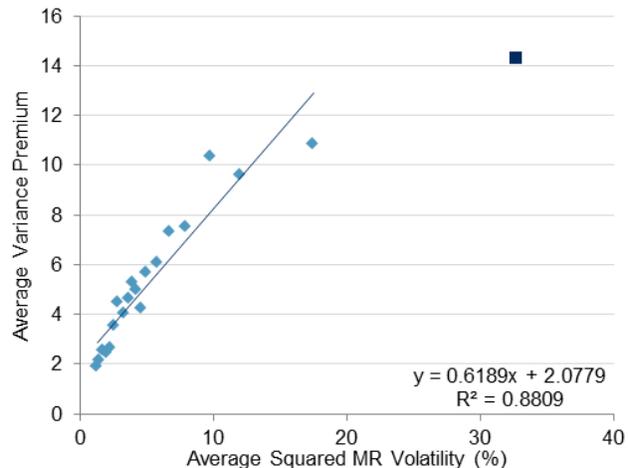
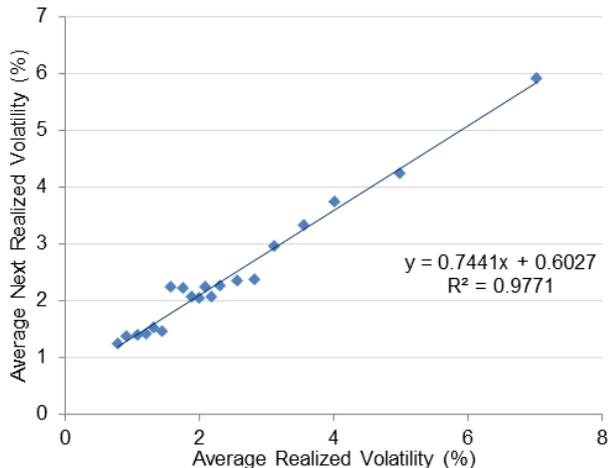
Sources: Hang Seng, S&P Dow Jones Indices LLC, and CBOE. Data from Jan. 2, 2014, to Sept. 29, 2017. VCR was calculated on each date and suffers from look-ahead bias prior to Sept. 29, 2017. Values for October 2017 were calculated assuming historical relationships were

unchanged. These values for VCR compared to changes in recent volatility in the Hong Kong Hang Seng Index over the corresponding 30-day period. Past performance is no guarantee of future results. Chart is provided for illustrative purposes.

S&P/JPX JGB VIX

Recent volatility in the [S&P 10-Year JGB Futures Total Return Index](#) has shown a strong tendency to mean revert; the goodness-of-fit (R-squared) is 0.98 for the linear relationship between average recent volatility and next realized volatility. A similarly strong R-squared value of 0.88 is observed in the linear relationship between the squared MR volatility and the VP; this compares to an R-squared figure of 0.65 for the linear relationship between MR volatility and the difference (between the Hang Seng Volatility Index and MR volatility).

Exhibit 28: MR and VP of the S&P/JPX JGB VIX



Source: S&P Dow Jones Indices LLC. Data from Feb. 13, 2008, to Sept. 29, 2017. Recent volatility and next realized volatility for the S&P 10-year JGB Futures Total Return were calculated on each day using closing price levels in the index from Jan. 11, 2008, to Oct. 31, 2017. Recent volatility and next realized volatility values were then evenly grouped into 20 buckets depending on the percentile rank of recent volatility across the whole period. The chart shows the average recent volatility and the average next realized volatility in each group. Past performance is no guarantee of future results. Chart is provided for illustrative purposes.

Source: S&P Dow Jones Indices LLC and CBOE. Data from Feb. 13, 2008, to Sept. 29, 2017. MR volatility was calculated using a constant $M=2\%$, $S=27\%$ for the period and suffers from look ahead bias. MR volatility was squared and we calculated the VP on each date. These values were then evenly grouped into 20 buckets depending on the percentile rank of recent volatility across the whole period. The chart shows the average MR volatility and the average VP in each group. The line of best fit was calculated using only the first 19 buckets; it ignored days when recent volatility is in the highest 5% of all readings. Past performance is no guarantee of future results. Chart is provided for illustrative purposes.

Exhibit 29: Performance of the VCR of the S&P/JPX JGB VIX



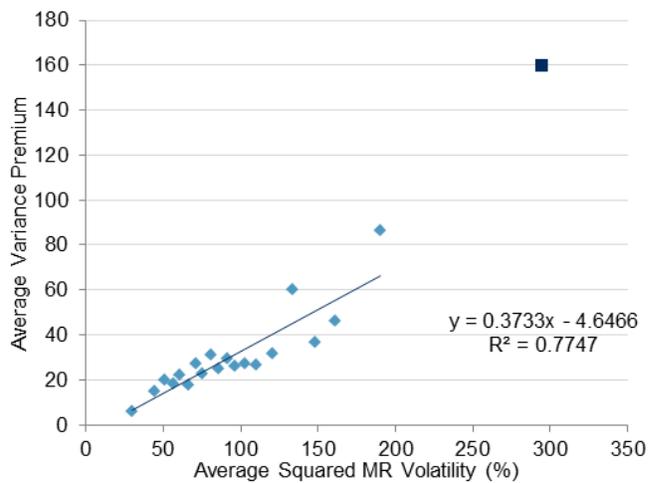
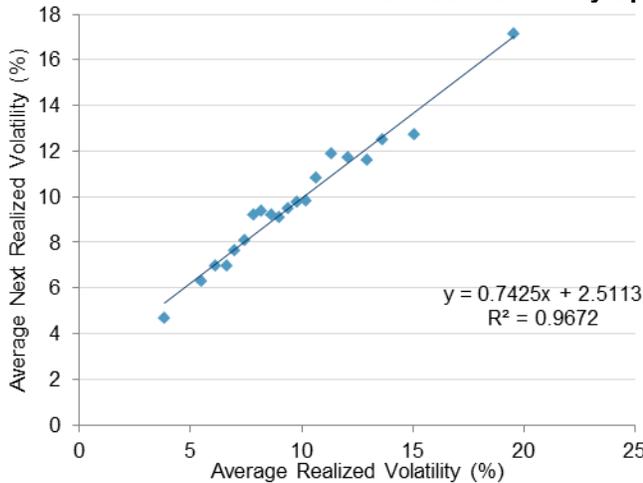
Source: S&P Dow Jones Indices LLC and CBOE. Data from Jan. 2, 2014, to Sept. 29, 2017. VCR was calculated on each date and suffers from look-ahead bias prior to Sept. 29, 2017. Values for October 2017 were calculated assuming historical relationships were unchanged. These values for VCR compared to changes in recent volatility in the S&P 10-year JGB Futures Total Return over the corresponding 30-day period. Past performance is no guarantee of future results. Chart is provided for illustrative purposes and reflects hypothetical historical

performance. Please see the Performance Disclosure at the end of this document for more information regarding the inherent limitations associated with back-tested performance.

CBOE/CME FX Euro Volatility

Recent volatility in the EUR/USD Currency Spot Rate has shown a strong tendency to mean revert; the goodness-of-fit (R-squared) is 0.97 for the linear relationship between average recent volatility and next realized volatility. A reasonably strong R-squared value of 0.78 is observed in the linear relationship between the squared MR volatility and the VP; this compares to an R-squared figure of 0.60 for the linear relationship between MR volatility and the difference (between the CBOE/CME FX Euro Volatility Index and MR volatility).

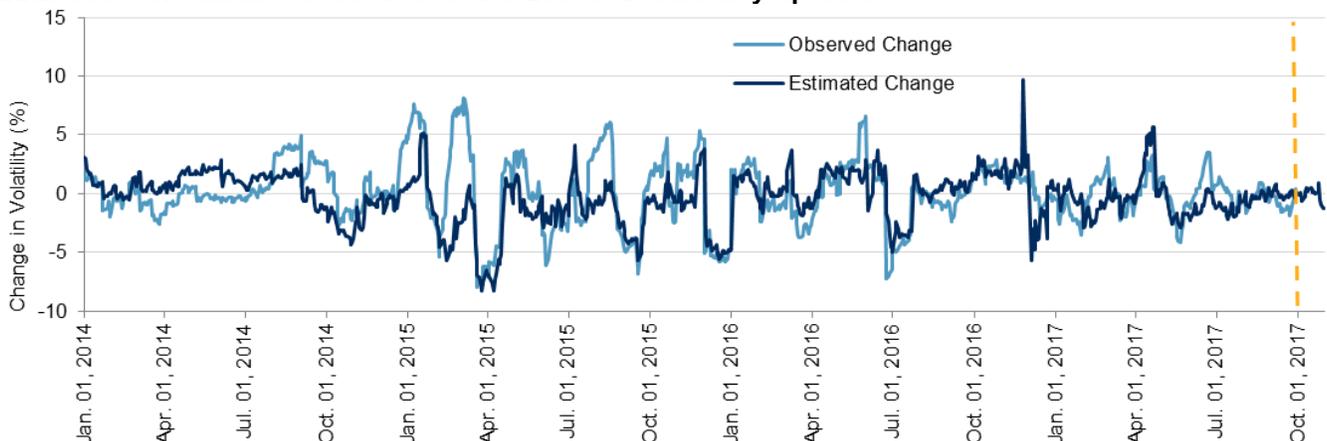
Exhibit 30: MR and VP of the EUR/USD Currency Spot Rate



Source: Reuters and S&P Dow Jones Indices LLC. Data from Jan. 7, 2008, to Sept. 29, 2017. Recent volatility and next realized volatility for the EUR/USD Spot Rate were calculated on each day using closing price levels in the index from Dec. 7, 2007, to Oct. 31, 2017. Recent volatility and next realized volatility values were then evenly grouped into 20 buckets depending on the percentile rank of recent volatility across the whole period. The chart shows the average recent volatility and the average next realized volatility in each group. Past performance is no guarantee of future results. Chart is provided for illustrative purposes

Source: Reuters, S&P Dow Jones Indices LLC and CBOE/CME. Data from Jan. 7, 2008, to Sept. 29, 2017. MR volatility was calculated using a constant M=10%, S=26% for the period and suffers from look-ahead bias. MR volatility was squared and we calculated the VP on each date. These values were then evenly grouped into 20 buckets depending on the percentile rank of recent volatility across the whole period. The chart shows the average MR volatility and the average VP in each group. The line of best fit was calculated using only the first 19 buckets; it ignored days when recent volatility is in the highest 5% of all readings. Past performance is no guarantee of future results. Chart is provided for illustrative purposes.

Exhibit 31: Performance of the VCR of the EUR/USD Currency Spot Rate



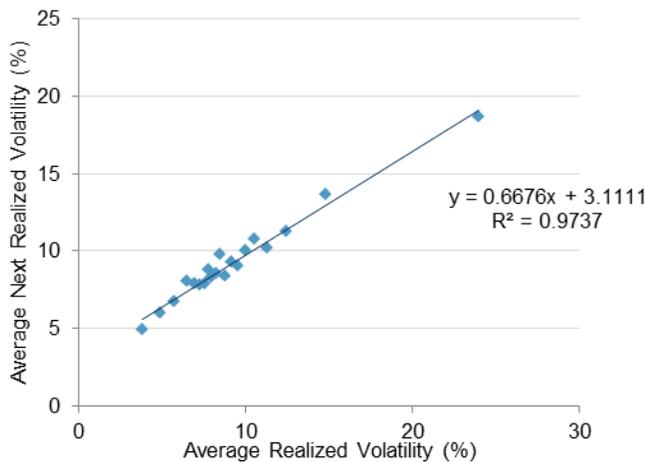
Source: Reuters and CBOE/CME. Data from Jan. 2, 2014, to Sept. 29, 2017. VCR is calculated on each date and suffers from look-ahead bias prior to Sept. 29, 2017. Values for October 2017 were calculated assuming historical relationships were unchanged. These values for

VCR compared to changes in recent volatility in the EUR/USD Spot Rate over the corresponding 30-day period. Past performance is no guarantee of future results. Chart is provided for illustrative purposes.

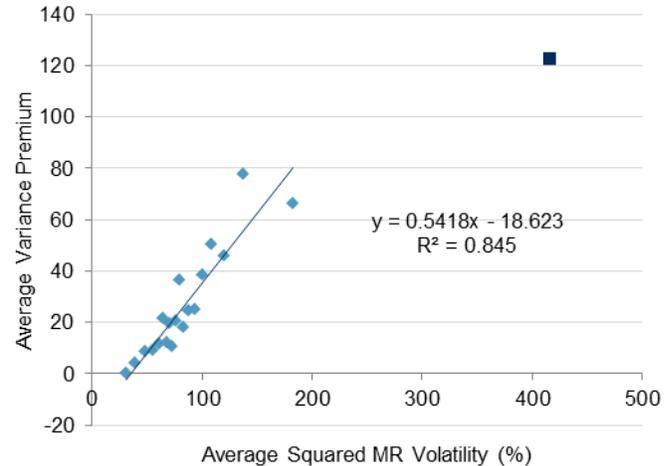
CBOE/CME FX GBP Volatility

Recent volatility in the GBP/USD Currency Spot Rate has shown a strong tendency to mean revert; the goodness-of-fit (R-squared) is 0.97 for the linear relationship between average recent volatility and next realized volatility. A reasonably strong R-squared value of 0.84 is observed in the linear relationship between the squared MR volatility and the VP; this compares to an R-squared figure of 0.79 for the linear relationship between MR volatility and the difference (between the CBOE/CME FX GBP Volatility Index and MR volatility).

Exhibit 32: MR and VP of the GBP/USD Currency Spot Rate



Source: Reuters and S&P Dow Jones Indices LLC. Data from Jan. 7, 2008, to Sept. 29, 2017. Recent volatility and next realized volatility for the GBP/USD Spot Rate were calculated on each day using closing price levels in the index from Dec. 7, 2007, to Oct. 31, 2017. Recent volatility and next realized volatility values were then evenly grouped into 20 buckets depending on the percentile rank of recent volatility across the whole period. The chart shows the average recent volatility and the average next realized volatility in each group. Past performance is no guarantee of future results. Chart is provided for illustrative purposes.



Source: Reuters, S&P Dow Jones Indices LLC and CBOE/CME. Data from Jan. 7, 2008, to Sept. 29, 2017. MR volatility was calculated using a constant $M=10\%$, $S=28\%$ for the period and suffers from look-ahead bias. MR volatility was squared and we calculated the VP on each date. These values were then evenly grouped into 20 buckets depending on the percentile rank of recent volatility across the whole period. The chart shows the average MR volatility and the average VP in each group. The line of best fit was calculated using only the first 19 buckets; it ignored days when recent volatility is in the highest 5% of all readings. Past performance is no guarantee of future results. Chart is provided for illustrative purposes.

Exhibit 33: Performance of the VCR of the GBP/USD Currency Spot Rate

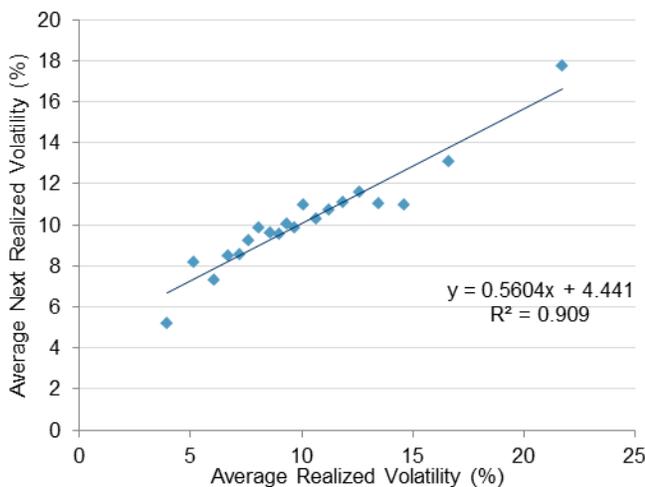


Source: Reuters and CBOE/CME. Data from Jan. 2, 2014, to Sept. 29, 2017. VCR is calculated on each date and suffers from look-ahead bias prior to Sept. 29, 2017. Values for October 2017 were calculated assuming historical relationships were unchanged. These values for VCR compared to changes in recent volatility in the GBP/USD Spot Rate over the corresponding 30-day period. Past performance is no guarantee of future results. Chart is provided for illustrative purposes.

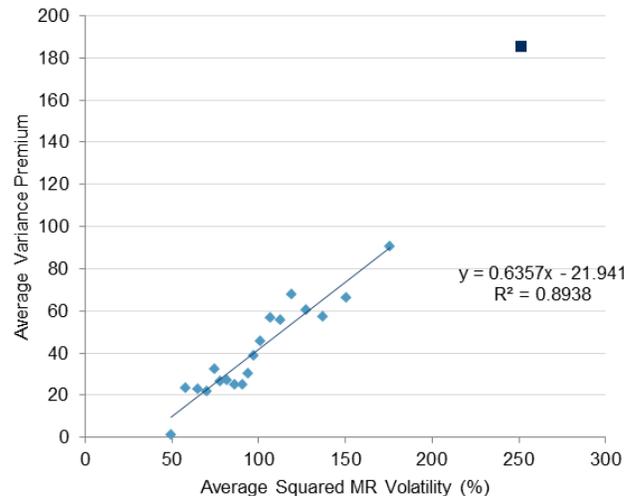
CBOE/CME FX Yen Volatility

Recent volatility in the Yen/USD Currency Spot Rate has shown a strong tendency to mean revert; the goodness-of-fit (R-squared) is 0.91 for the linear relationship between average recent volatility and next realized volatility. The same R-squared value (0.90) is observed in the linear relationship between the squared MR volatility and the VP; this compares to an R-squared figure of 0.75 for the linear relationship between MR volatility and the difference (between the CBOE/CME FX Yen Volatility Index and MR volatility).

Exhibit 34: MR and VP of the Yen/USD Currency Spot Rate

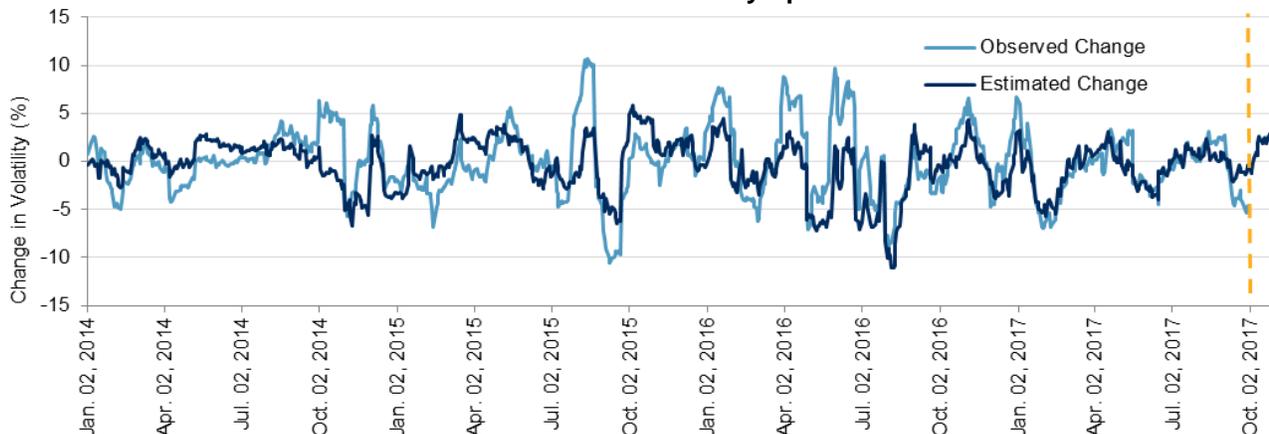


Source: Reuters and S&P Dow Jones Indices LLC. Data from Jan. 7, 2008, to Sept. 29, 2017. Recent volatility and next realized volatility for the Yen/USD spot rate were calculated on each day using closing price levels in the index from Dec. 3, 2007, to Oct. 31, 2017. Recent volatility and next realized volatility values were then evenly grouped into 20 buckets depending on the percentile rank of recent volatility across the whole period. The chart shows the average recent volatility and the average next realized volatility in each group. Past performance is no guarantee of future results. Chart is provided for illustrative purposes.



Source: Reuters, S&P Dow Jones Indices LLC and CBOE/CME. Data from Jan. 7, 2008, to Sept. 29, 2017. MR volatility was calculated using a constant M=10%, S=51% for the period and suffers from look-ahead bias. MR volatility was squared and we calculated the VP on each date. These values were then evenly grouped into 20 buckets depending on the percentile rank of recent volatility across the whole period. The chart shows the average MR volatility and the average VP in each group. The line of best fit was calculated using only the first 19 buckets; it ignored days when recent volatility is in the highest 5% of all readings. Past performance is no guarantee of future results. Chart is provided for illustrative purposes.

Exhibit 35: Performance of the VCR of the Yen/USD Currency Spot Rate

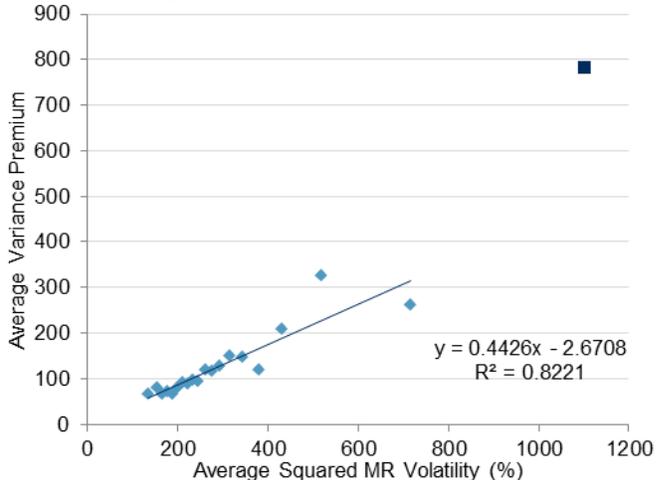
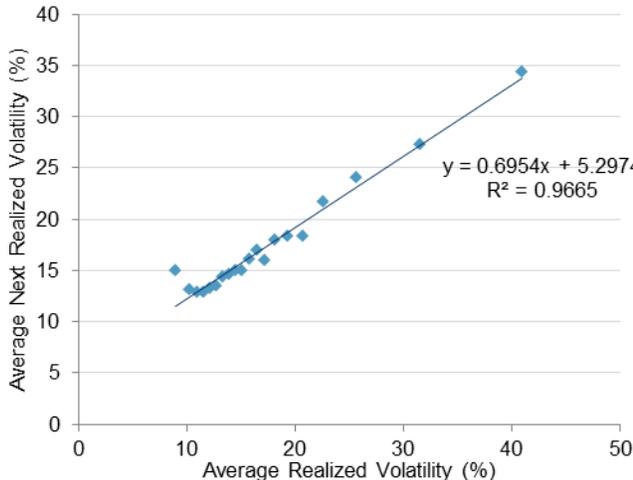


Source: Reuters and CBOE and CME. Data from Jan. 2, 2014, to Sept. 29, 2017. VCR is calculated on each date and suffers from look-ahead bias prior to Sept. 29, 2017. Values for October 2017 were calculated assuming historical relationships were unchanged. These values for VCR compared to changes in recent volatility in the Yen/USD spot rate over the corresponding 30-day period. Past performance is no guarantee of future results. Chart is provided for illustrative purposes.

CBOE Gold ETF Volatility

Recent volatility in the Gold London Pm Fix rate has shown a strong tendency to mean revert; the goodness-of-fit (R-squared) is 0.97 for the linear relationship between average recent volatility and next realized volatility. The R-squared value of 0.82 is observed in the linear relationship between the squared MR volatility and the VP; this compares to an R-squared figure of 0.67 for the linear relationship between MR volatility and the difference (between the CBOE Gold ETF Volatility Index and MR volatility).

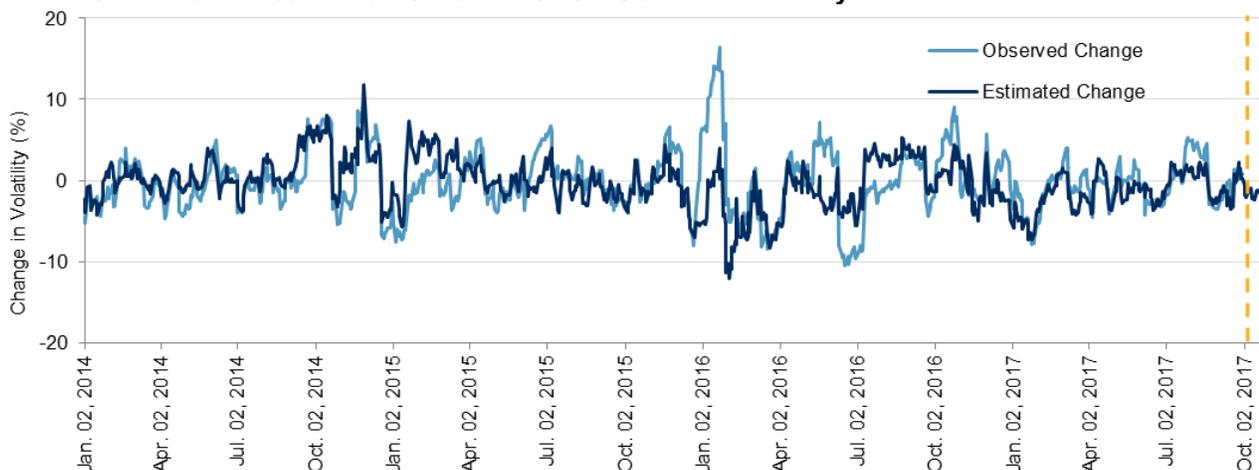
Exhibit 36: Mean Reversion and VP of the CBOE Gold ETF Volatility



Source: ICE and S&P Dow Jones Indices LLC. Data from June 3, 2008, and Sept. 29, 2017, Recent volatility and next realized volatility for the GOLD London PM Fix were calculated on each day using closing price levels in the index from May 2, 2008, to Oct. 31, 2017. Recent volatility and next realized volatility values were then evenly grouped into 20 buckets depending on the percentile rank of recent volatility across the whole period. The chart shows the average recent volatility and the average next realized volatility in each group. Past performance is no guarantee of future results. Chart is provided for illustrative purposes.

Source: ICE, S&P Dow Jones Indices LLC and CBOE. Data from June 3, 2008, to Sept. 29, 2017. MR volatility was calculated using a constant M=17%, S=33% and suffers from look-ahead bias. MR volatility was squared and we calculated the VP on each date. These values were then evenly grouped into 20 buckets depending on the percentile rank of recent volatility across the whole period. The chart shows the average MR volatility and the average VP in each group. The line of best fit was calculated using only the first 19 buckets; it ignored days when recent volatility is in the highest 5% of all readings. Past performance is no guarantee of future results. Chart is provided for illustrative purposes.

Exhibit 37: Performance of the VCR of the CBOE Gold ETF Volatility

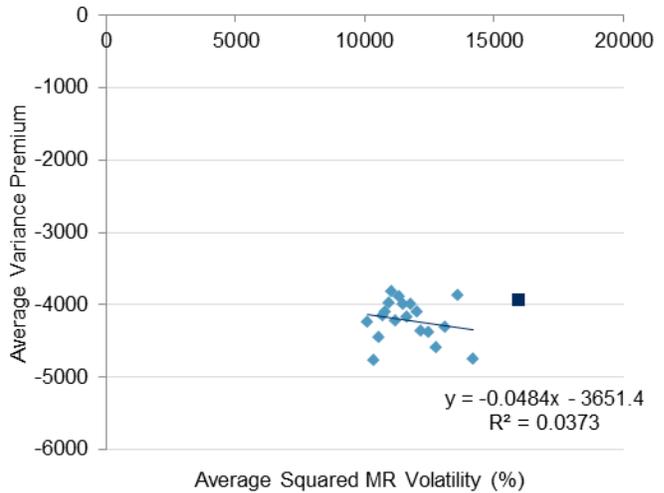
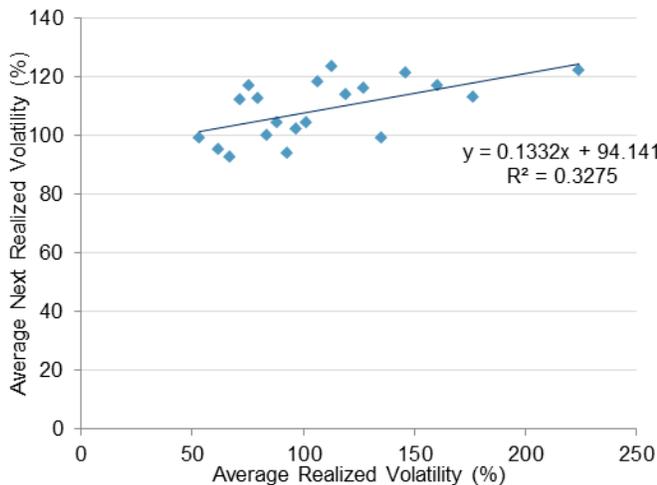


Source: ICE/CBOE. Data from Jan. 2, 2014, to Sept. 29, 2017. VCR was calculated on each date and suffers from look-ahead bias prior to Sept. 29, 2017. Values for October 2017 were calculated assuming historical relationships were unchanged. These values for VCR compared to changes in recent volatility in the GOLD London PM Fix over the corresponding 30-day period. Past performance is no guarantee of future results. Chart is provided for illustrative purposes.

CBOE VIX of VIX

Recent volatility in VIX has shown a slight tendency to mean revert; the goodness-of-fit (R-squared) is 0.33 for the linear relationship between average recent volatility and next realized volatility. An extremely small R-squared value of 0.04 is observed in the linear relationship between the squared MR volatility and the VP; this compares to an R-squared figure of 0.25 for the linear relationship between MR volatility and the difference (between the VIX and MR volatility).

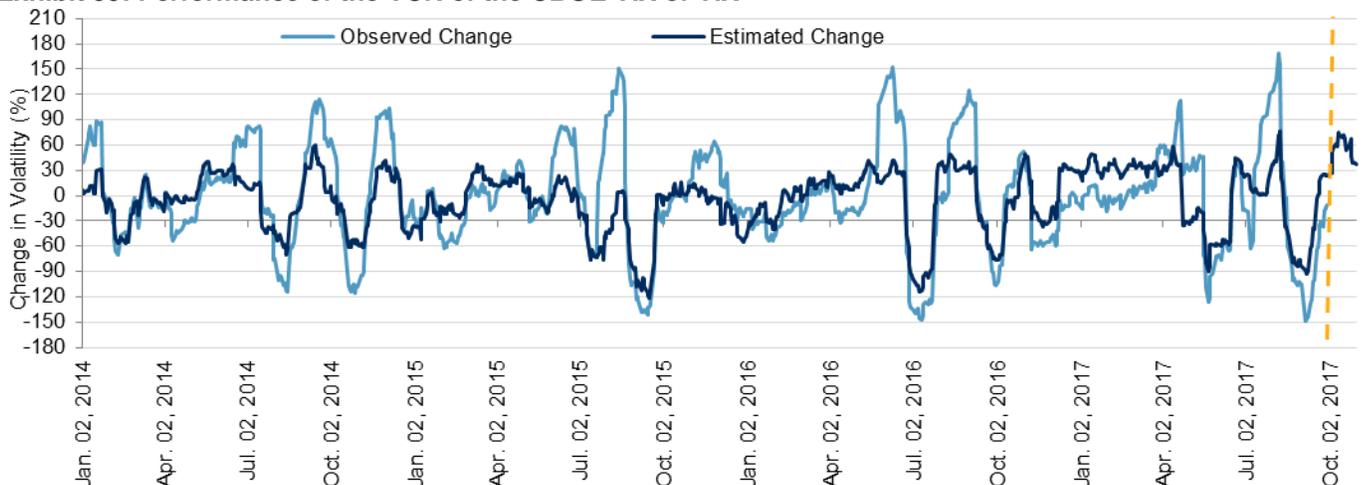
Exhibit 38: MR and VP of the CBOE VIX of VIX



Source: S&P Dow Jones Indices LLC and CBOE. Data from March 6, 2006, to Sept. 29, 2017. Recent volatility and next realized volatility for VIX were calculated on each day using closing price levels in the index from Feb. 2, 2006, to Oct. 31, 2017. Recent volatility and next realized volatility values were then evenly grouped into 20 buckets depending on the percentile rank of recent volatility across the whole period. The chart shows the average recent volatility and the average next realized volatility in each group. Past performance is no guarantee of future results. Chart is provided for illustrative purposes.

Source: S&P Dow Jones Indices LLC and CBOE. Data from March 6, 2006, to Sept. 29, 2017. MR volatility was calculated using a constant M=109%, S=85% for the period and suffers from look-ahead bias. MR volatility was squared and we calculated the VP on each date. These values were then evenly grouped into 20 buckets depending on the percentile rank of recent volatility across the whole period. The chart shows the average MR volatility and the average VP in each group. The line of best fit was calculated using only the first 19 buckets; it ignored days when recent volatility is in the highest 5% of all readings. Past performance is no guarantee of future results. Chart is provided for illustrative purposes.

Exhibit 39: Performance of the VCR of the CBOE VIX of VIX



Source: S&P Dow Jones Indices LLC and CBOE. Data from Jan. 2, 2014, to Sept. 29, 2017. VCR is calculated on each date and suffers from look-ahead bias prior to Sept. 29, 2017. Values for October 2017 were calculated assuming historical relationships were unchanged. These values for VCR compared to changes in recent volatility in the VIX over the corresponding 30-day period. Past performance is no guarantee of future results. Chart is provided for illustrative purposes.

S&P DJI RESEARCH CONTRIBUTORS		
Charles Mounts	Global Head	charles.mounts@spglobal.com
Jake Vukelic	Business Manager	jake.vukelic@spglobal.com
GLOBAL RESEARCH & DESIGN		
AMERICAS		
Aye M. Soe, CFA	Americas Head	aye.soe@spglobal.com
Dennis Badlyans	Associate Director	dennis.badlyans@spglobal.com
Phillip Brzenk, CFA	Director	phillip.brzenk@spglobal.com
Smita Chirputkar	Director	smita.chirputkar@spglobal.com
Rachel Du	Senior Analyst	rachel.du@spglobal.com
Bill Hao	Director	wenli.hao@spglobal.com
Qing Li	Associate Director	qing.li@spglobal.com
Berlinda Liu, CFA	Director	berlinda.liu@spglobal.com
Ryan Poirier, FRM	Senior Analyst	ryan.poirier@spglobal.com
Maria Sanchez	Associate Director	maria.sanchez@spglobal.com
Kelly Tang, CFA	Director	kelly.tang@spglobal.com
Peter Tsui	Director	peter.tsui@spglobal.com
Hong Xie, CFA	Director	hong.xie@spglobal.com
APAC		
Priscilla Luk	APAC Head	priscilla.luk@spglobal.com
Utkarsh Agrawal, CFA	Associate Director	utkarsh.agrawal@spglobal.com
Liyu Zeng, CFA	Director	liyu.zeng@spglobal.com
Akash Jain	Associate Director	akash.jain@spglobal.com
EMEA		
Sunjiv Mainie, CFA, CQF	EMEA Head	sunjiv.mainie@spglobal.com
Leonardo Cabrer, PhD	Senior Analyst	leonardo.cabrer@spglobal.com
Andrew Innes	Senior Analyst	andrew.innes@spglobal.com
INDEX INVESTMENT STRATEGY		
Craig J. Lazzara, CFA	Global Head	craig.lazzara@spglobal.com
Fei Mei Chan	Director	feimei.chan@spglobal.com
Tim Edwards, PhD	Senior Director	tim.edwards@spglobal.com
Anu R. Ganti, CFA	Director	anu.ganti@spglobal.com
Hamish Preston	Senior Associate	hamish.preston@spglobal.com
Howard Silverblatt	Senior Industry Analyst	howard.silverblatt@spglobal.com

PERFORMANCE DISCLOSURE

The S&P 500 was launched on March 4, 1957. The S&P/JPX JGB VIX was launched on October 2, 2015. All information presented prior to an index's Launch Date is hypothetical (back-tested), not actual performance. The back-test calculations are based on the same methodology that was in effect on the index Launch Date. Complete index methodology details are available at www.spdji.com.

S&P Dow Jones Indices defines various dates to assist our clients in providing transparency. The First Value Date is the first day for which there is a calculated value (either live or back-tested) for a given index. The Base Date is the date at which the Index is set at a fixed value for calculation purposes. The Launch Date designates the date upon which the values of an index are first considered live: index values provided for any date or time period prior to the index's Launch Date are considered back-tested. S&P Dow Jones Indices defines the Launch Date as the date by which the values of an index are known to have been released to the public, for example via the company's public website or its datafeed to external parties. For Dow Jones-branded indices introduced prior to May 31, 2013, the Launch Date (which prior to May 31, 2013, was termed "Date of introduction") is set at a date upon which no further changes were permitted to be made to the index methodology, but that may have been prior to the Index's public release date.

Past performance of the Index is not an indication of future results. Prospective application of the methodology used to construct the Index may not result in performance commensurate with the back-test returns shown. The back-test period does not necessarily correspond to the entire available history of the Index. Please refer to the methodology paper for the Index, available at www.spdji.com for more details about the index, including the manner in which it is rebalanced, the timing of such rebalancing, criteria for additions and deletions, as well as all index calculations.

Another limitation of using back-tested information is that the back-tested calculation is generally prepared with the benefit of hindsight. Back-tested information reflects the application of the index methodology and selection of index constituents in hindsight. No hypothetical record can completely account for the impact of financial risk in actual trading. For example, there are numerous factors related to the equities, fixed income, or commodities markets in general which cannot be, and have not been accounted for in the preparation of the index information set forth, all of which can affect actual performance.

The Index returns shown do not represent the results of actual trading of investable assets/securities. S&P Dow Jones Indices LLC maintains the Index and calculates the Index levels and performance shown or discussed, but does not manage actual assets. Index returns do not reflect payment of any sales charges or fees an investor may pay to purchase the securities underlying the Index or investment funds that are intended to track the performance of the Index. The imposition of these fees and charges would cause actual and back-tested performance of the securities/fund to be lower than the Index performance shown. As a simple example, if an index returned 10% on a US \$100,000 investment for a 12-month period (or US \$10,000) and an actual asset-based fee of 1.5% was imposed at the end of the period on the investment plus accrued interest (or US \$1,650), the net return would be 8.35% (or US \$8,350) for the year. Over a three year period, an annual 1.5% fee taken at year end with an assumed 10% return per year would result in a cumulative gross return of 33.10%, a total fee of US \$5,375, and a cumulative net return of 27.2% (or US \$27,200).

GENERAL DISCLAIMER

Copyright © 2017 by S&P Dow Jones Indices LLC, a part of S&P Global. All rights reserved. Standard & Poor's®, S&P 500® and S&P® are registered trademarks of Standard & Poor's Financial Services LLC ("S&P"), a subsidiary of S&P Global. Dow Jones® is a registered trademark of Dow Jones Trademark Holdings LLC ("Dow Jones"). Trademarks have been licensed to S&P Dow Jones Indices LLC. Redistribution, reproduction and/or photocopying in whole or in part are prohibited without written permission. This document does not constitute an offer of services in jurisdictions where S&P Dow Jones Indices LLC, Dow Jones, S&P or their respective affiliates (collectively "S&P Dow Jones Indices") do not have the necessary licenses. All information provided by S&P Dow Jones Indices is impersonal and not tailored to the needs of any person, entity or group of persons. S&P Dow Jones Indices receives compensation in connection with licensing its indices to third parties. Past performance of an index is not a guarantee of future results.

It is not possible to invest directly in an index. Exposure to an asset class represented by an index is available through investable instruments based on that index. S&P Dow Jones Indices does not sponsor, endorse, sell, promote or manage any investment fund or other investment vehicle that is offered by third parties and that seeks to provide an investment return based on the performance of any index. S&P Dow Jones Indices makes no assurance that investment products based on the index will accurately track index performance or provide positive investment returns. S&P Dow Jones Indices LLC is not an investment advisor, and S&P Dow Jones Indices makes no representation regarding the advisability of investing in any such investment fund or other investment vehicle. A decision to invest in any such investment fund or other investment vehicle should not be made in reliance on any of the statements set forth in this document. Prospective investors are advised to make an investment in any such fund or other vehicle only after carefully considering the risks associated with investing in such funds, as detailed in an offering memorandum or similar document that is prepared by or on behalf of the issuer of the investment fund or other vehicle. Inclusion of a security within an index is not a recommendation by S&P Dow Jones Indices to buy, sell, or hold such security, nor is it considered to be investment advice.

These materials have been prepared solely for informational purposes based upon information generally available to the public and from sources believed to be reliable. No content contained in these materials (including index data, ratings, credit-related analyses and data, research, valuations, model, software or other application or output therefrom) or any part thereof (Content) may be modified, reverse-engineered, reproduced or distributed in any form or by any means, or stored in a database or retrieval system, without the prior written permission of S&P Dow Jones Indices. The Content shall not be used for any unlawful or unauthorized purposes. S&P Dow Jones Indices and its third-party data providers and licensors (collectively "S&P Dow Jones Indices Parties") do not guarantee the accuracy, completeness, timeliness or availability of the Content. S&P Dow Jones Indices Parties are not responsible for any errors or omissions, regardless of the cause, for the results obtained from the use of the Content. THE CONTENT IS PROVIDED ON AN "AS IS" BASIS. S&P DOW JONES INDICES PARTIES DISCLAIM ANY AND ALL EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, ANY WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE OR USE, FREEDOM FROM BUGS, SOFTWARE ERRORS OR DEFECTS, THAT THE CONTENT'S FUNCTIONING WILL BE UNINTERRUPTED OR THAT THE CONTENT WILL OPERATE WITH ANY SOFTWARE OR HARDWARE CONFIGURATION. In no event shall S&P Dow Jones Indices Parties be liable to any party for any direct, indirect, incidental, exemplary, compensatory, punitive, special or consequential damages, costs, expenses, legal fees, or losses (including, without limitation, lost income or lost profits and opportunity costs) in connection with any use of the Content even if advised of the possibility of such damages.

S&P Dow Jones Indices keeps certain activities of its business units separate from each other in order to preserve the independence and objectivity of their respective activities. As a result, certain business units of S&P Dow Jones Indices may have information that is not available to other business units. S&P Dow Jones Indices has established policies and procedures to maintain the confidentiality of certain non-public information received in connection with each analytical process.

In addition, S&P Dow Jones Indices provides a wide range of services to, or relating to, many organizations, including issuers of securities, investment advisers, broker-dealers, investment banks, other financial institutions and financial intermediaries, and accordingly may receive fees or other economic benefits from those organizations, including organizations whose securities or services they may recommend, rate, include in model portfolios, evaluate or otherwise address.