1ST QUARTER ▪ 2019

INVESTMENT PUBLICATIONS HIGHLIGHTS

First quarter's edition summarizes five articles on the varying role volatility plays in the investment process:

- the first article considers the likelihood that future equity premiums will be positive after accounting for equity market volatility across multiple investment horizons;
- the second argues the level of risk present in portfolios is misrepresented when investors use volatility measures that do not align with investment horizons;
- the third assesses the impact of investors' asymmetric risk preferences on long-term equity returns by leveraging the options market to quantify the volatility risk premium;
- the fourth highlights different equity risk management strategies and finds that their fundamental appeal is their ability to reduce tail risks; and
- the last article analyzes price volatility before, during, and after 40 well-known historical financial bubbles and finds that heightened volatility is not a reliable indicator of market crashes.

VOLATILITY LESSONS

Eugene F. Fama and Kenneth R. French, *Financial Analysts Journal*, vol 74, no 3 (Third Quarter 2018): 42–53

The authors studied the extent to which equity market volatility impacts equity, value, and small-cap premiums over medium- and long-term time horizons. They find a high probability of earning negative premiums across time horizons and caution that investors should not base expected returns on the average historical return.

To study volatility's impact on the equity market premium, the authors collected a dataset of monthly stock market returns from a market cap–weighted portfolio of NYSE, AMEX, and NASDAQ stocks from July 1963 to December 2016, along with monthly Treasury bill returns for the same period. The authors analyzed return horizons ranging from one month to 30 years. Because available sample sizes shrink as the horizon period is extended, for periods of 12 months or more, the authors simulated performance of more than 100,000 iterations from randomly selected paired monthly equity and Treasury bill returns. They employed two methods in their simulation; the first used the dataset as given, whereas the second introduced an uncertainty component by incorporating the standard error of the dataset. The authors also studied volatility's impact on value and small-cap premiums using the same time periods and simulation methods.

The average monthly equity market premium from the dataset was 0.51% per month, or roughly 6% per annum. But the standard deviation of the monthly equity premium was a substantial 4.42%, which the authors argue raises questions about whether the future equity premiums will be positive. Likewise, the average monthly value and small-cap premiums were 0.29% and 0.27%, respectively, but their standard deviations were a sizable 2.19% and 2.82%, respectively.

Using simulated data over three- and five-year periods, negative equity market premiums occurred 28.5% and 23.4% of the time, respectively. Even at longer time horizons, negative results occurred 15.6% of the time over ten-year periods, and 7.9% of the time over 20-year periods. The value premium was negative 23.2% and 17.1% of the time over three- and five-year horizons, respectively, and the small-cap premium was negative 34.1% and 29.8% over those same periods, respectively. Using the second simulation method—which employs the sampling's standard error estimates to create a wider range of observations—the chances of earning a negative premium increase.

Positive skewness in the simulated distributions increased as the return horizon grew. Thus, the authors explained that good outcomes became more likely, and more extreme, than bad outcomes. The authors note this characteristic is positive, but caution investors regarding volatility nonetheless. Given the high likelihood of earning a negative premium across return horizons simply by chance, they argue that equity return expectations should not be solely based on average realized returns.

TIME DIVERSIFICATION REDUX

Michael Aked and Amie Ko, Research Affiliates LLC, August 2017

The authors examine investment risk and find that conventional measures of volatility may be misleading. They show that measures of the annualized volatility of US equity returns vary as the length of the holding period used to calculate volatility is changed. This finding leads the authors to suggest that investors should rely on risk measures that match their individual investment horizon to get a more accurate representation of the risk in their portfolios.

Time diversification—the notion that time diversifies risk—suggests that the volatility of risky assets decreases over longer holding periods. But there are statistical limitations associated with measuring volatility over longer holding periods. As a result, it has become standard practice for the financial industry to use ever-shorter time horizons to measure volatility. The authors believe that this trend has major implications for assessing portfolio investment risk.

By measuring the annualized standard deviation of US equities from 1871 to 2016 over one-month, three-month, one-year, five-year, and ten-year holdings periods, the authors found that changing the holding period used to calculate this volatility metric resulted in varying measures of risk. Volatility steadily increased for holding periods of one month to one year—peaking at one year—and then steadily declined for holding periods of one year to ten years. This pattern held true for 15 different global asset classes and a sample of 86 US mutual funds with at least a 50-year track record.

The authors posit that volatility changes with the holding period because of two well-known financial concepts: momentum and mean reversion. The increase in volatility

over holding periods of one-month to one-year is evidence of a positive correlation between monthly returns, also known as the momentum effect. Whereas mean reversion, or the negative correlation between monthly returns over longer periods, is driving the decrease in volatility over holding periods of one year to ten years. The authors acknowledge that, over time, uncertainty overwhelms the effects of mean reversion and volatility eventually increases again over very long holding periods.

These findings lead the authors to conclude that time diversification is helpful, at least over certain horizons; specifically, after return momentum has faded (around one year) and before the investment horizon becomes so long that uncertainty takes hold. Furthermore, assessing risk by calculating annualized volatility using daily or monthly holding periods may not be an accurate representation of the actual risk present in many investment portfolios. Instead, long-term investors should consider risk in the context of their investment horizon and rely on measures of annualized volatility that align with their longer holding periods.

EMBRACING DOWNSIDE RISK

Roni Israelov, Lar N. Nielson, and Daniel Villalon, *The Journal of Alternative Investments*, Winter 2017: 59–67

Investors have asymmetric risk preferences, which explains the existence of a *volatility risk premium*—compensation for bearing downside risk. The authors leverage the options market to quantify the volatility risk premium to assess the impact of investors' asymmetric risk preferences on US equity returns. They find that the volatility risk premium is the main source of long-term equity returns.

It is known that investors have asymmetric risk preferences: they covet upside participation but have a weak tolerance for downside risk. This risk aversion means that acquiring downside protection comes at a cost (commonly referred to as the volatility risk premium). Investors regularly use the options market to hedge their exposure to downside risk. Within the options market, the volatility risk premium equals the difference between the option's implied volatility and the underlying asset's realized volatility. A larger volatility risk premium reflects more extreme asymmetric risk preferences.

The authors use the options market to measure the impact of asymmetric risk preferences on US equity returns. To achieve this goal, they split the S&P 500 Index into upside and downside components. The upside component—equity exposure with a long call position—has limited downside risk and unlimited upside participation. The downside component—equity exposure with a covered call position—has full downside participation but a capped upside. If investors have asymmetric risk preferences, then investors demanding downside protection will pay their counterparty (i.e., investors accepting downside risk) a volatility risk premium.

Looking at the historical data from 1986 to 2014, the authors observe that the downside component, which accepts risk, outperformed the upside component, which hedges risk, by 500 basis points. In fact, the authors found that roughly 80% of the S&P 500's cumulative returns over this period can be attributed to the downside component. Meanwhile, the upside component contributes very little reward, but

accounts for almost half of the S&P 500's total risk. In short, the S&P 500 is risky, but compensation for bearing this risk is asymmetric. This reflects investor preferences, making the large volatility risk premium payment for downside protection unpalatable in the long run.

The authors ran the same analysis for bonds, gold, commodity futures, and credit, and the results support their findings from the original analysis; downside risk exposure accounts for the majority of returns. The authors conclude that if the downside risk for a given asset is too much for certain investors; a better option is to reduce the size of their allocation to that asset. Less risk-averse investors may consider embracing downside risk, as the large premium associated with taking this position will likely lead to superior long-term outcomes.

THE STATE OF RISK MANAGEMENT

Corey Hoffstein and Justin Sibears, Newfound Research, August 20, 2018

The authors explore the efficacy of equity risk management strategies from a risk-adjusted return perspective. They argue that although many of these strategies limit upside performance, many have proven to be effective at mitigating tail risks. As a result, several strategies have generated superior Sharpe ratios relative to a buy and hold strategy.

The researchers examined eight risk-mitigation strategies and discovered that six strategies outperformed a relevant benchmark from December 1997 to July 2018 on a risk-adjusted basis. The strategies leveraged well-known techniques/tools, such as a 60/40 mix of stocks and bonds, trend-following measures, and factor-based equity allocations. The two strategies that underperformed—an option put protection and an option collar strategy—experienced larger drawdowns and more volatility.

The authors also compared these strategies over shorter time periods. As might be expected, results varied over rolling one-year relative returns. But the strategies with the widest dispersion of returns (e.g., trend following) offered the most downside protection. Said another way, the strategies with the solid downside protection had greater levels of short-term volatility.

In general, risk management strategies underperformed during market upswings and outperformed during times of crisis. For example, during the dot-com bust and the global financial crisis, risk management strategies outperformed the S&P 500 81.2% of the time, but only 19.8% of the time during the 17 years in which the S&P 500 posted positive returns. This demonstrates the insurance-like quality of these strategies; investors using these strategies forego relative upside performance to reduce downside capture during a market decline.

The authors also sought to determine if diversifying across the risk management strategies added value. Interestingly, an equal-weighted blend of all the observed strategies (excluding collar/protective put strategies) outperformed the individual returns of seven out of eight risk-management strategies. This demonstrated a key benefit when implementing risk management strategies. Investors could opt for a diversified blend of strategies rather than try to hand pick the best strategy.

In their conclusion, the authors propose an introspective question: "Why bother with risk management in the first place?" To answer this question, two Monte Carlo simulations were conducted to assess the performance of the S&P 500 over a 30-year time horizon; one simulation was conducted without a risk management overlay and the other with an equal-weighted blend of risk management strategies. The results showed that the distribution of returns was much narrower when risk management strategies were employed. This narrowing, or reduction in tail risk, is the fundamental appeal of risk management.

CAN WE USE VOLATILITY TO DIAGNOSE FINANCIAL BUBBLES? LESSONS FROM 40 HISTORICAL BUBBLES

Didier Sornette, Peter Cauwels, and Georgi Smilyanov, Swiss Finance Institute, *Research Paper Series*, no 17-27

The authors observe price volatility before, during, and after financial bubbles to explore whether volatility is a leading indicator of market crashes. They consider 40 well-known historical bubbles and find little evidence to suggest heightened volatility precedes asset bubble peaks; rather, periods of low volatility are more common before a crash. The authors conclude that price volatility is not a useful indicator for identifying the end of asset-price bubbles.

Financial bubbles typically occur when financial assets trade at prices substantially higher than their fundamental value. When a bubble pops, prices fall in dramatic fashion and often lead to wider market panic, such as at the end of the late-1990s dot-com bubble, the US housing bubble that preceded the 2008 Great Recession, and the Dutch "tulip mania" in the early 17th century. Academics have rigorously studied financial bubbles, yet investors are still in search of a practical model for detecting when bubbles will burst. One common perception is that the price volatility of an asset experiencing a bubble spikes before the price crashes.

To test whether this theory holds up in practice, the authors inspected price volatility before, during, and after 40 well-known asset bubbles that subsequently crashed. The authors define volatility as the annualized rolling 20-day standard deviation of the daily change in asset prices. The 40 bubbles (occurring between 1929 and 2011) covered multiple asset classes, including developed and emerging markets equities, individual stocks, ETFs, currencies, and commodities. The authors plotted asset price and price volatility data for the 200-day window around asset price peaks to determine whether volatility tends to spike as an asset's price reaches its peak. Each bubble is then categorized as either "fearful" (volatility and prices grow together) or "fearless" (volatility does not grow with price).

Contrary to the generally accepted paradigm, the authors find that price volatility exhibits no consistent relationship with asset prices in the lead up to a crash. Take, for instance, the striking similarities between the Malaysian and Philippine bubbles; both bubbles had common drivers that developed over the same timeframe and burst simultaneously. Yet, volatility declined right before the market crash in Malaysia, while it increased prior to the one in the Philippines. In fact, across all 40 bubbles included in the analysis, the authors observed more cases of "fearless" bubbles than "fearful" bubbles. In other words, price volatility was more likely to be low as opposed to high before a bubble burst. These findings indicate that volatility may not be a reliable tool for identifying maturing asset bubbles that are poised to crash. ■