THE TIME-VARIANT RISK OF EQUITY INDICES

by CataMetrics

CataMetrics®
MANAGEMENT, LLC

Printed in the United States of America.

Contents

An equity strategy that protects investors from high volatilities should have a good chance of providing improved returns with a lower level of risk. However, the level of risk in equities is not constant over time and this variability complicates the construction of investment portfolios. Specifically, for investment solutions that span multiple periods, not considering the time-variant nature of equity risk may lead to intractable discrepancies between forecast and realized results.

This paper contributes to the discussion of the problem of time-variant equity risk by presenting some empirical observations relating to the S&P 500 index as manifested through the behavior of the exchange-traded fund (ETF) SPY. We demonstrate that the same observations apply generally to developed-market equities. While this paper can be read as a stand-alone observation of equity-index return characteristics, it is primarily meant as a technical companion to a forthcoming piece on CataMetrics Tactical Strategies which will explore more comprehensively the problem of equity risk over time in portfolio construction.

Portfolio construction in equity markets with time-variant risks presents a problem to investors. Namely, assuming that the distribution of investment returns repeats itself from one time period to another allows for elegant and simple single-period analytics but the price of this convenience is that actual investment outcomes may be difficult to reconcile with the expected outcomes based on invariant return distributions. In other words, for investment solutions that span multiple periods, not considering the time-variant nature of equity risks may lead to large differences between forecast and realized returns.

In this paper we use the term 'time diversification' as an analogue to the concept of cross-sectional diversification: Just as a few large and risky allocations can dominate the overall riskiness of a portfolio in any one time period, we show that the riskiness over time of a portfolio can be driven largely by high risk during a limited number of time periods. We use equity-index volatility and returns to illustrate our conclusions though, of course, the conclusions carry over to investment portfolios in general.[1]

Finally, a note of caution to avoid confusion: A reader familiar with life-cycle products will recognize the term 'time diversification' as commonly used to describe how the riskiness of equites varies as a function of an investment's time horizon. This is a different, albeit related, usage of 'time diversification', one that we will not address.[2]

[1] Ashwin Alankar and Myron Scholes of Janus Capital Group, in a publication of the Investment Insights Series named "Adaptive Asset Allocation Refocusing Portfolio Management Toward Investor End Goals", neatly illustrate the impact that volatile time periods can have on the long-term outcomes of investment portfolios.

[2] How the choice of investment time horizon affects the risk of equities is a somewhat contentious topic with important implications not only for the meaningfulness of life-cycle products but also for the tools and rationales used by financial advisors to determine the investment objectives and suitability of investments for their clients.

Two informative articles on this other usage of the term time diversification have been published by Donald G. Bennyhoff of Vanguard and Jack Duval of Duval Asset Management. Bennyhoff, under the header of Vanguard Investment Counseling & Research, published a carefully balanced and comprehensive review-article, "The Diversification and Horizon-Based Asset Allocation", about the arguments used to support and disprove the notion that a longer time horizon lowers the riskiness of equities. In the spring of 2006, Duval published "The Myth of Time Diversification: Analysis, Application, and Incorrect Account Forms", in which he argues that believing that equity risks decreases with longer time horizons can lead investment advisors to assign their clients inappropriate investment objectives.

[3] Until May 10, 2013, the indexation rules of S&P Dow Jones Indices, SPDJI, were that the opening prices for the SPX had to be calculated using the opening prices of the index's constituent stocks at 09:30:00. If an insufficient number of stocks traded directly at the open, the rule stipulated that the previous closing value would be used as the opening price. A consequence of the rule was that on many days the previous night's closing level was carried forward. As of May 13, 2013, SPDJI delayed the opening price-take by one second to 09:30:01, and the problem was eliminated.

This lack of reliable historical opening level for equity indices is a problem ▸

Measuring Risk, Returns, and Risk-adjusted Returns – and why the SPY

Our analysis is based on measuring realized risks and returns, and risk-adjusted returns, over short time periods. While measuring returns is straightforward, the measurement of risk is more complicated theoretically and the time periods for which we are able to analyze short-term volatility are limited because of a lack of granular-enough data.

To measure the changing volatility of equities over time, we measure realized volatility over short, non-overlapping time windows. Our methodology employs short-term volatility estimates and in order to make the most of available data for short time periods, or, in other words, to have the most efficient mechanism for estimating market volatility, we use as inputs the opening, high, low, and closing prices from each day in our estimation period.

In an ideal world we would have used the opening, high, low, and closing prices of actual equity indices but even though the closing levels are readily available for a broad range of equity indices, there is surprisingly considerably less availability of meaningful opening index levels.[3]

The solution to this data problem is to use index ETFs instead because the stock-exchange data for the opening, high, low, and closing prices is more comprehensive than for many of the ETFs' underlying indices.

The longest times series of equity index data with high quality observations of the open, high, low, and closing prices turns out to be for the SPY, the underlying index of which is the S&P 500, which also was the first US ETF. The data that we have used for our analysis runs from February 12, 1993, up to and including April 8, 2016.[4]

The formula that we use is a derivative of what is generally referred to as the Garman-Klass Yang-Zhang extension.[5, 6]

We choose to base our analysis on non-overlapping ten-day windows for each of which we calculate the volatility (using the open, high, low, and close for each of the ten days) and return of the SPY. For summary measures, we also divide observations of the ten-day windows into sextiles defined by the volatility estimates.

The choice for the length of the estimation periods and the use of sextiles is somewhat arbitrary but the conclusions would not change if we used a different number of days for the estimation periods or some other quantile measure. Sextiles, moreover, make for elegant diagrams.

▸ since leaving out the overnight movements in an index would make it difficult to use the opening, high, low, closing price methodology in a meaningful way.

Interestingly enough, this rule change does not seem to have been implemented for the ten GICS indices that make up the SPX (with the Bloomberg tickers S5COND, S5CONS, S5ENRS, S5FINL, S5HLTH, S5INDU, S5INFT, S5MATR, S5TELS, and S5UTIL, respectively) since the Bloomberg data for these ten indices for the period as of May 13, 2013 continue to show a very high frequency of observations where the opening price equals the previous closing price. Note that this paper was written before the real-estate sector was split off from the financials sector at the end of August 2016.

4 The first day of trading for the SPY was January 29, 1993, but, to make sure that the market for this new instrument was operating properly, we excluded the first ten days of trading from our analysis.

For reference, as of the second day of trading of the SPY, February 1, 1993, until May 10, 2013, there were 5,107 trading days on the NYSE and during this time the SPY's opening price equaled its previous closing price 133 times, or about 3% of the time. The corresponding numbers for the SPX are 3,162, or 62% of the time.

5 There is a nice review of different volatility estimators by Colin Bennet and Miguel A. Gil of Groupo Santander from February 3, 2012, titled "Measuring Historical Volatility".

6 We make minor modifications to the standard formula in addition to dividend-adjusting the prices of the ETFs. The modification entails replacing the annualization factor in Garman-Klass Yang-Zhang with a more precise measure that takes into account calendar days rather than just an assumed number of trading days per year. The purpose of adjusting the observed market prices of the ETFs is to avoid the small volatility spikes that otherwise would occur for estimation periods during which an ETF goes ex-dividend. The price adjustments entail using dividend re-invested prices as well as expressing these prices relative to the return of Fed funds, our proxy for the risk-free rate.

Observation #1: Risk is Distributed Unevenly Over Time

Our first observation is that the risk of the SPY is distributed unevenly over time.

Diagram 1 is a familiar depiction of how the volatility of the ETF has varied over time: There is the unmistakable 100% + spike from October 2008, the 'death of volatility' during 2004 – 2006, the second largest spike which followed 9/11, and other high volatility periods amongst them the third highest volatility spike in early September, 2015, following the gyrations of the Chinese stock market.[7]

Diagram 1: 10-day Annualized Volatility of the SPY with an Average of 16.3%

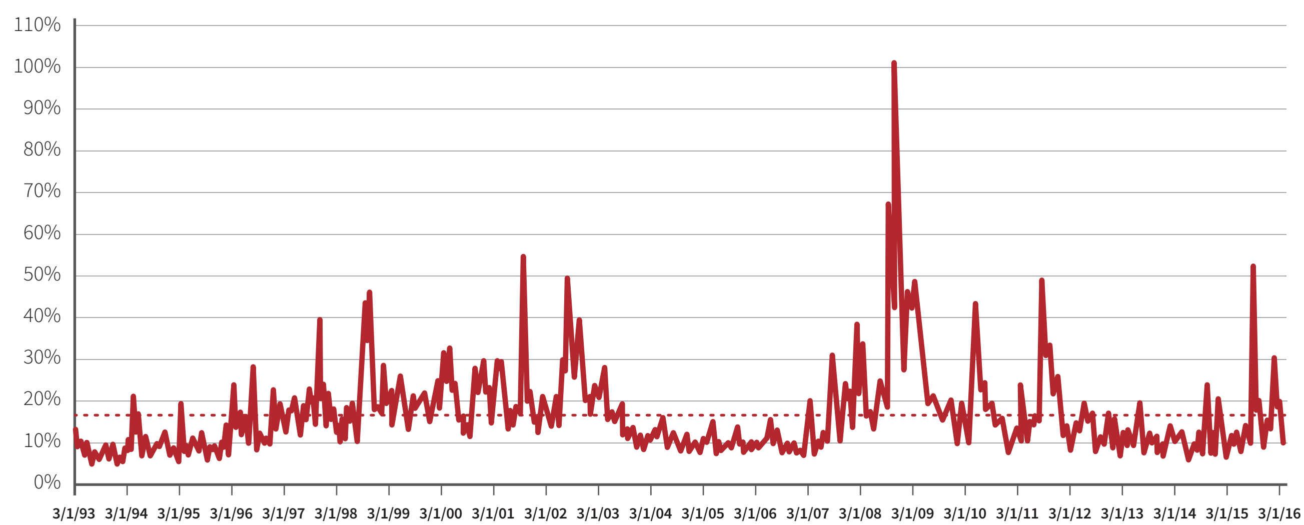

For reference, and this is represented by the dotted line in the diagram, the average volatility of the 583 ten-day periods, which make up the red line in the diagram, is 16.3%.

Diagram 1 shows that the periods with above average volatility have contributed considerably more to the total variability of the SPY than have the periods with below average volatility. Just looking at the red graph, however, does not give much guidance to quantifying 'considerably more'.

[7] The period runs from the ten-day period ending with March 1, 1993, to the ten-day period ending with April 8, 2016.

To depict how massively skewed the influence of the volatile periods is on the total variability of the SPY, we categorize the 583 non-overlapping ten-day periods into sextiles based on the realized volatility of the periods.[8] For each of the sextiles we then calculate the variance (not the volatility[9]) of its constituent observations, and the resulting distribution of the total variance is shown in *Diagram 2* below.

Diagram 2: Distribution of Variance Mass across Sextiles for the SPY

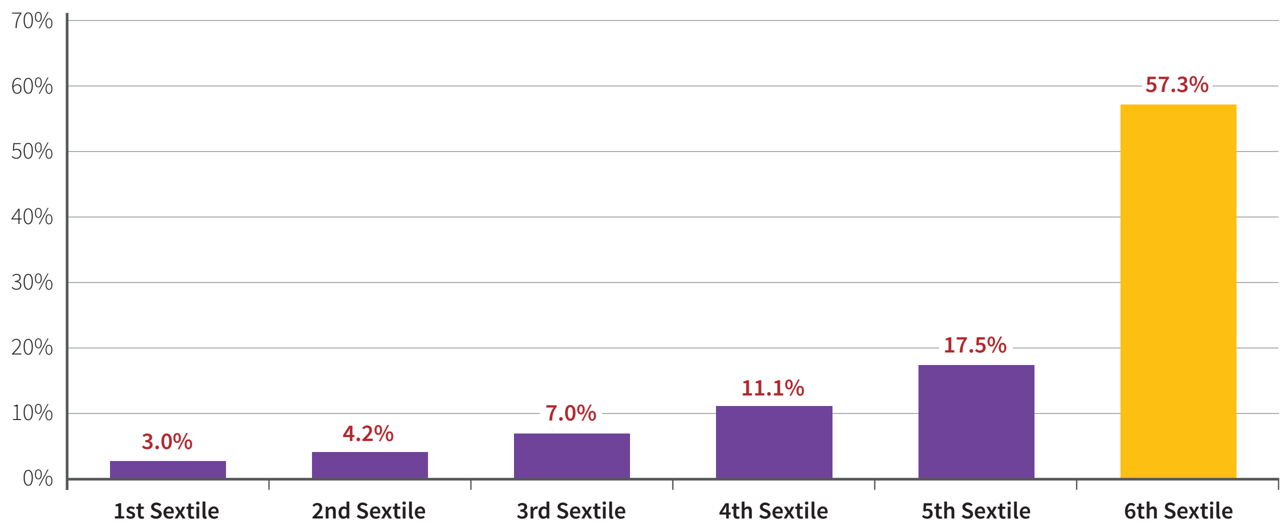

It is striking that more than half of the variance mass, 57.3% in fact, is attributable to the butterscotch-coloured bar representing the sixth sextile. Sextiles six and five together represent 74.7% of the variance mass and since the lower volatility limit of the 5th sextile, 17.3%, is just above the average volatility of 16.3% as given in *Diagram 1,* it is safe to state that more than three-quarters of the variance mass is attributable to periods which have above average volatility.

'Considerably more', in other words, means that in the aggregate, the periods with above average volatility contributed to the variability of the SPY more than three times than what the periods with below average volatility contributed.

What happens to returns in volatile periods thus has a major impact on the performance of the SPY over time.

[8] The upper limits of the first five sextiles in terms of annualized ten-day volatilities are 8.98%, 10.74%, 14.01%, 17.32%, and 22.50%, respectively, with the sixth sextile encompassing all observations with volatilities higher than the upper limit of the fifth sextile. Note that the median volatility observation is equal to the upper limit of the third sextile, i.e. 14.01%, which is lower than the average of 16.32%. There are 583 non-overlapping ten-day periods with 98 observations in the first sextile and 97 observations in each of the other five sextiles.

[9] We use variance rather than volatility because variance measures are additive, so it is simple to assign variance mass, in this case the sum of the variance generated by the individual ten-day periods belonging to each sextile, across the sextiles. It is the variance mass of each sextile that determines how much of the total variability of the SPY over the whole time period is attributable to a particular sextile.

Observation #2: Risk is Bad for Equity Returns

Since we already have bucketed all non-overlapping ten-day periods into sextiles, it is easy to determine the average return of the observations in each sextile and the results are shown in *Diagram 3.*

It is striking how the average return falls with the level of the volatility. The first four of the sextiles have positive average returns, the fifth sextile has an average return of all but zero, and the sixth sextile has a significant average return of -1.37%. It certainly seems that calm markets tend to be conducive to positive equity returns while turbulent markets, on average, are associated with negative equity-market returns.

Diagram 3: Average Returns across Sextiles for the SPY

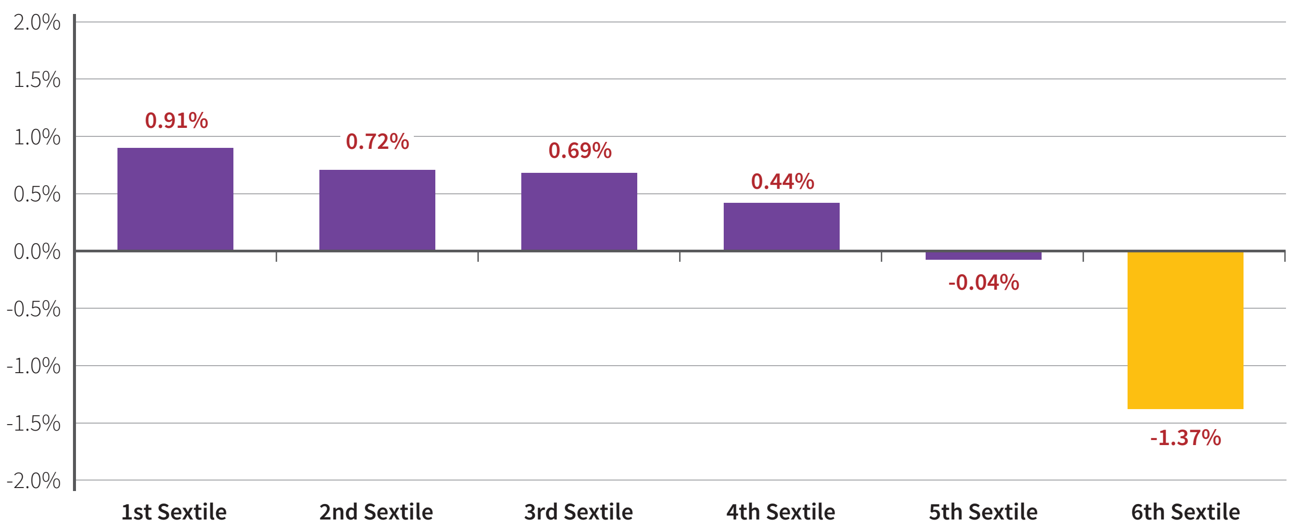

While striking, the bars in *Diagram 3* give us no information about the dispersion of returns within each sextile, meaning that the bars do not tell us the extent to which the returns within each sextile are concentrated or spread around the sextile's average return. For example, without a measure describing the spread of returns within each sextile, we have no indication whether the strong positive returns in the first sextile are due to a handful of positive outliers or if the negative average return of the sixth sextile is due to a couple of cataclysmic events in an otherwise stable environment.

The detailed ten-day observations that make up the data in *Diagrams 2* and *3* are displayed in *Diagram 4* below.

The axes in *Diagram 4* show the period volatilities and returns of the non-overlapping 10-day intervals meaning that we have not annualized any of the output parameters. The lavender-coloured circles show outcomes that make up sextiles one through five and the butterscotch-coloured dots represent the observations of the sixth sextile.

The scatter plot shows clearly the contrasting outcomes between sextiles one and six. The outcomes that define the first sextile, that is, outcomes with a period volatility not exceeding about 1.8%, are clearly concentrated above the horizontal axis. Conversely, the dots representing the observations of the sixth sextile show the proclivity of highly volatile periods to generate negative returns. While there are some materially positive outcomes in this sextile, the negative outcomes dominate. Furthermore, the negative average outcome of the observations of the sixth sextile is not caused by just a small number of extreme events—clearly, high volatility tends to be bad for returns.

Diagram 4: Period Returns and Volatilities of 10-day Intervals for the SPY

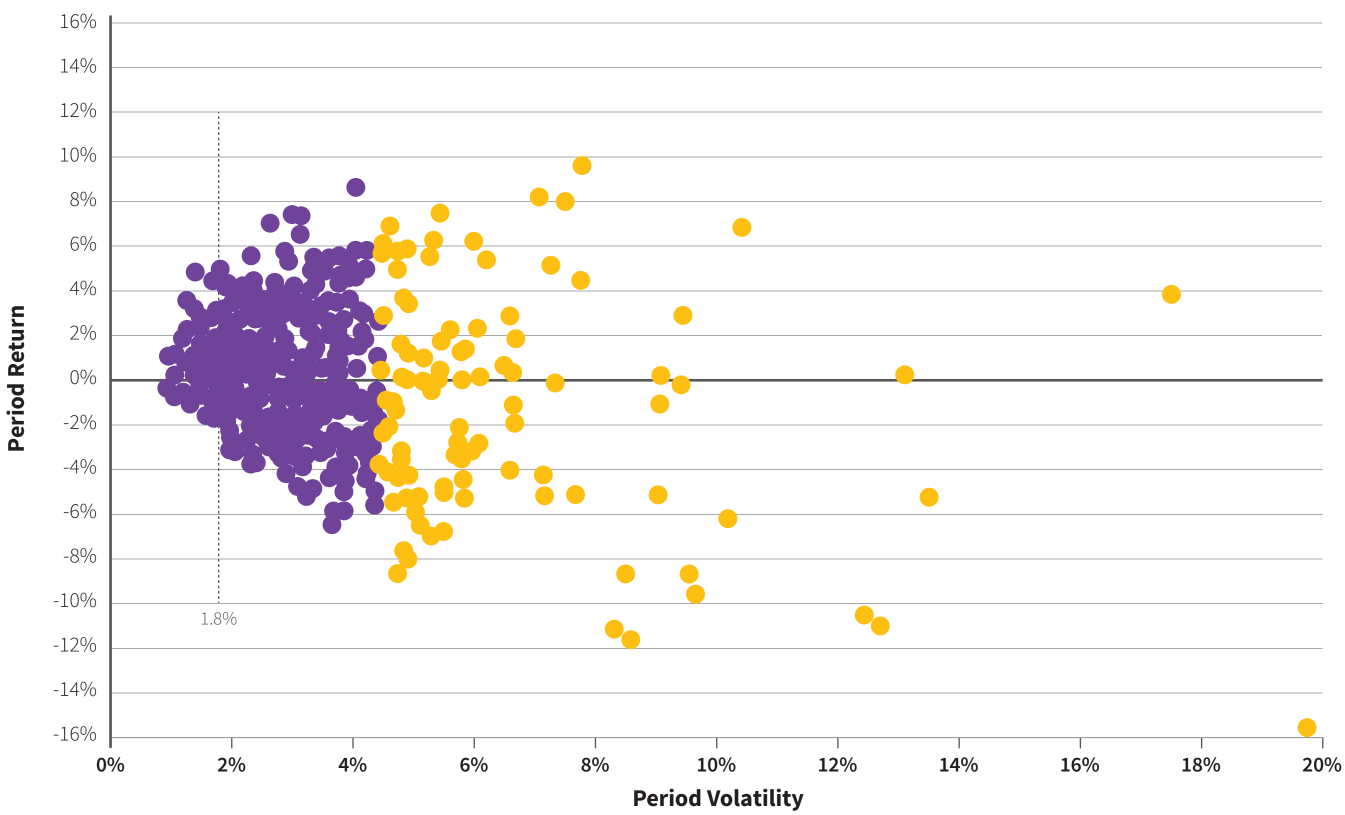

An intuitive way to summarize the information shown in *Diagram 2* through *Diagram 4* is to estimate the total information ratios (TIR)[10] for each of the sextiles. The notion of the TIR is simply to express the long-term ratio of return to risk for each sextile.

The resulting TIR are shown in *Diagram 5* and the conclusion is that over the lifetime of the SPY, which in this case means from mid-February 1993 to mid-April 2016, market risk has not been an investor's friend.

Diagram 5: Total Information Ratios across Sextiles for the SPY

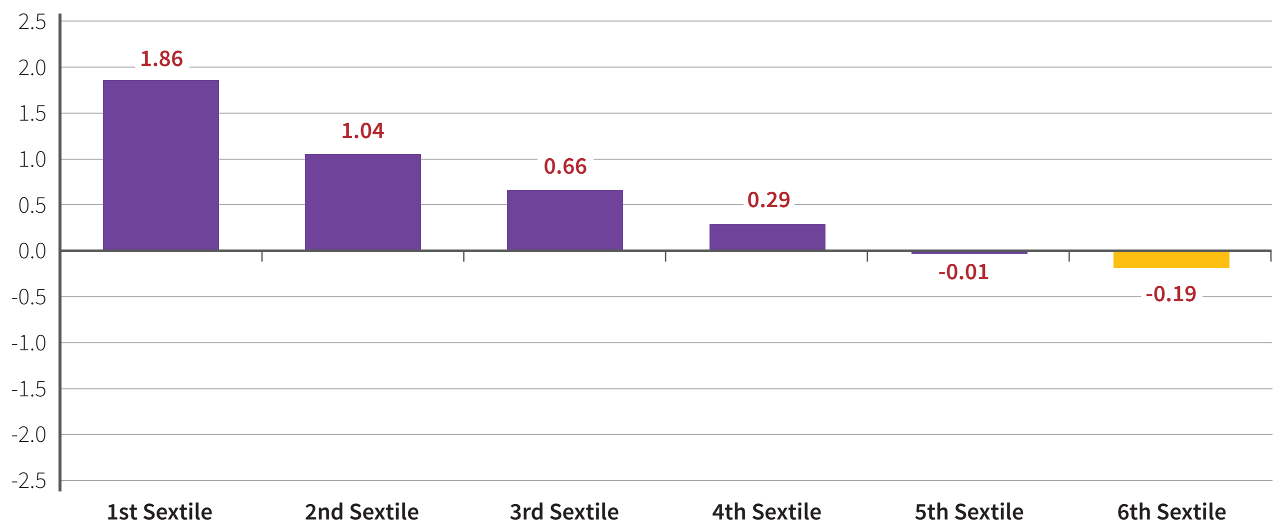

[10] We define the total information ratio to be the ratio between the accumulated returns in each sextile divided by the square root of the variance mass that the observations of each sextile have generated; note that in this paper we express returns as excess returns over the risk-free rate (for which we use the Fed funds rate). This measure gives the total return per unit of standard deviation for each sextile and is a more informative measure than the average of the information ratios of the individual observations that make up each sextile.

Risk and Return as Distributions

Another way to explore the effects on the returns of the SPY of the observations in the sixth sextile is to create frequency diagrams with both observed and fitted distributions. Since the underlying data is the same that we used for the sextile diagrams, the conclusions, of course, will be the same albeit that the presentation may be more familiar.

Diagram 6 shows the 583 return observations organized in buckets of size 0.005 for the logs of the returns. For each bucket, the observations for the first five sextiles make up the base of each bucket and the observations of the sixth sextile are on top.

There are not enough observations to create textbook-like smooth histograms but there are enough of them to illustrate clearly that the historical return distribution's left-hand tail is extended, and that the bulk of the tail is comprised of observations from the sixth sextile.

Diagram 6: Distribution of Observed Returns

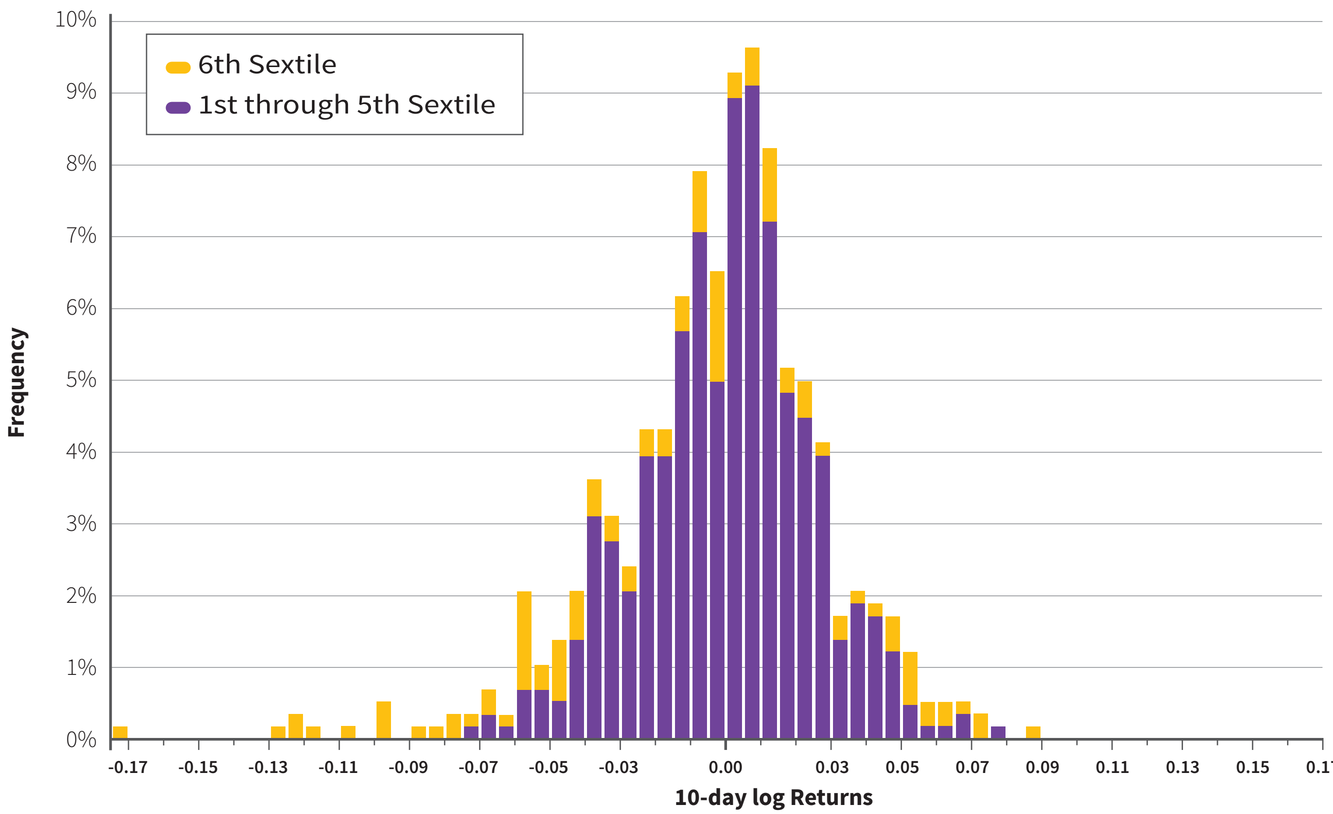

The fitted return distributions shown in *Diagram 7* are an attempt to quantify approximately the effect of the high-volatility observations on the overall historical return distribution. We use the mean and the variance of the observations from the first five sextiles, as shown in *Diagram 6,* to define the lavender-coloured fitted lognormal distribution in *Diagram 7.* In other words, this distribution shows what the historical distribution would have looked like in the absence of the observations that make up the sixth sextile—we can call this the 'partial' distribution.

Similarly, we use the parameters based on all observation returns to define the fitted return distribution for whole time-series of historical returns—we can call this the 'full' distribution—which is represented by the burgundy graph in *Diagram 7.*

The difference between the two distributions in *Diagram 7* is thus an illustration of the effect of the observations in the sixth sextile, meaning that we can now quantify, approximately, what we meant when we concluded that high volatility is bad for returns.

Diagram 7: Fitted Return Distributions

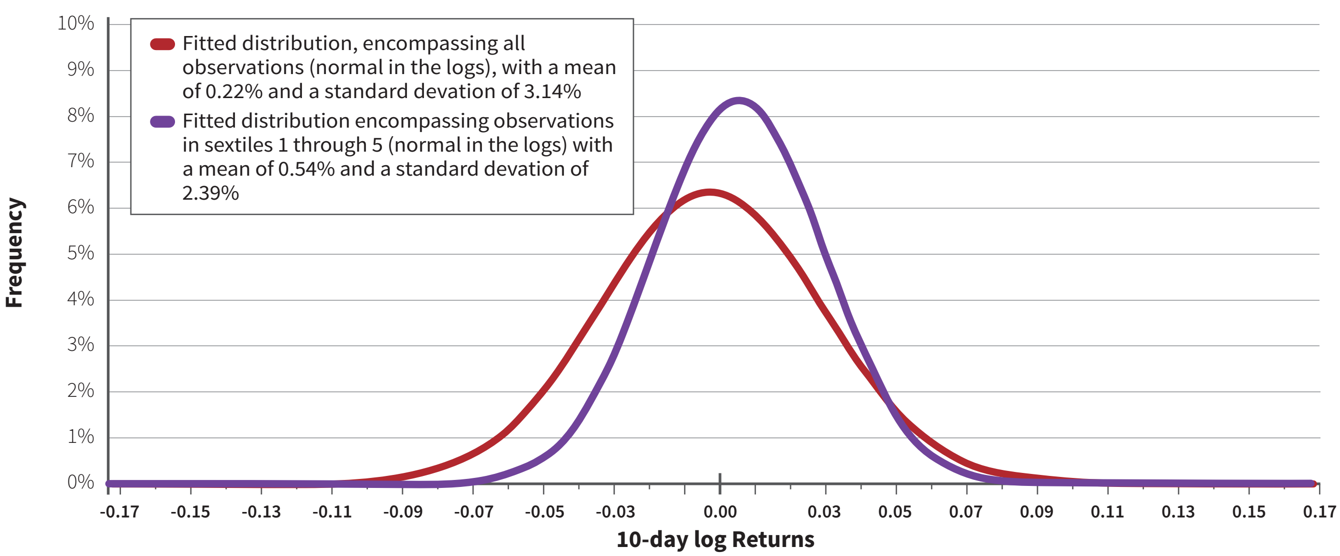

Since the mean for the partial and full distribution are 0.54% and 0.22%, respectively, adding the sixth sextile to the partial distribution means that, historically, the overall return was decreased by about three fifths. Likewise, an additional cost of going from the partial to the full distribution is an increase of the standard deviation from 2.39% to 3.14% for a ten-day period, which is an increase of about three tenths.

An International Comparison

The observations that we have made using data for the SPY carry over to other markets. To demonstrate this insight, we have used US market data for six other of the older and well-established equity ETFs: MDY for US mid cap, EWG for Germany, EWU for Britain, EWJ for Japan, EWH for Hong Kong, and EWS for Singapore.

The time period for this comparison starts with November 11, 1997, and runs to April 8, 2016, which is the end date for the analysis of the SPY.[11] For this comparison, we use exactly the same non-overlapping ten-day time periods as we do for the SPY except that the comparison starts in 1997 rather than in 1993, see *Diagram 9* through *Diagram 14.*

The comparison is made in terms of the total information ratio after each ETF's returns have been segmented into sextiles, just as we did with the SPY. For the sake of completeness and ease of comparison, we also show the outcomes for the same time, shorter, time period for the SPY in *Diagram 8* below.

[11] The choice of the time period for the comparison is governed by the first date as of which all six ETFs in the comparison have traded every day on the US exchanges; prior to November 11, 1997, every now and then, there were dates on which one or several of the ETFs did not trade. Because of illiquidity, we also excluded three other ETFs: EWA for Australia, EWC for Canada, and EWW for Mexico. Had we wanted to include these three ETFs, the start date for the comparison would have had to be in 2003, but we chose to go for a longer time period, with three fewer ETFs, with more observations covering both the fallout from the collapse of Long-Term Capital Management in 1998 and the bursting of the tech bubble in 2000.

Diagram 8: Total Information Ratios across Sextiles for the SPY

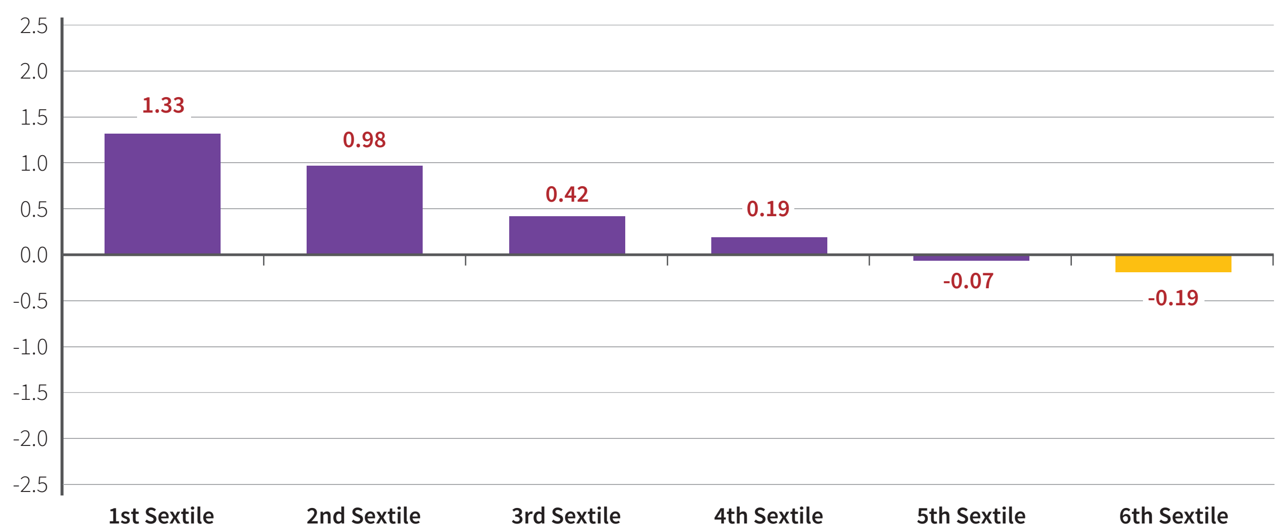

Diagram 9: Total Information Ratios across Sextiles for the MDY

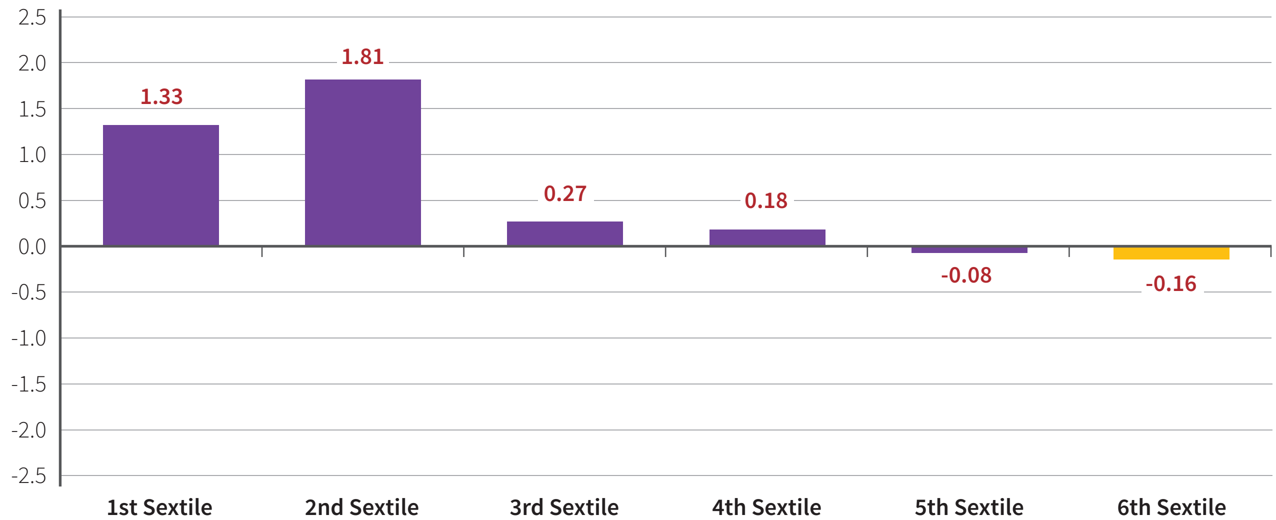

Diagram 10: Total Information Ratios across Sextiles for the EWU

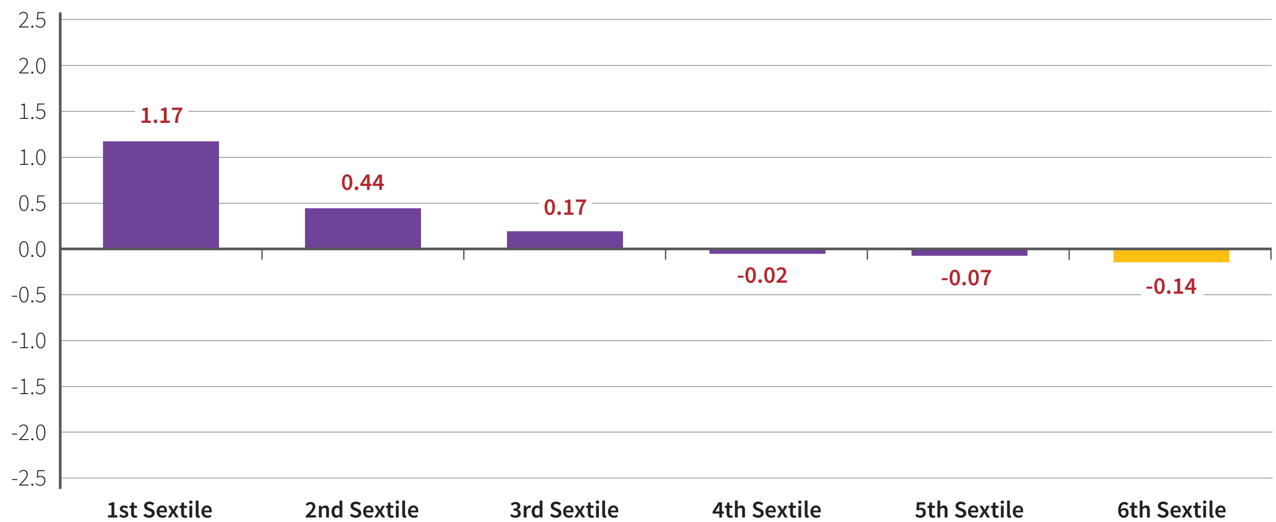

Diagram 11: Total Information Ratios across Sextiles for the EWG

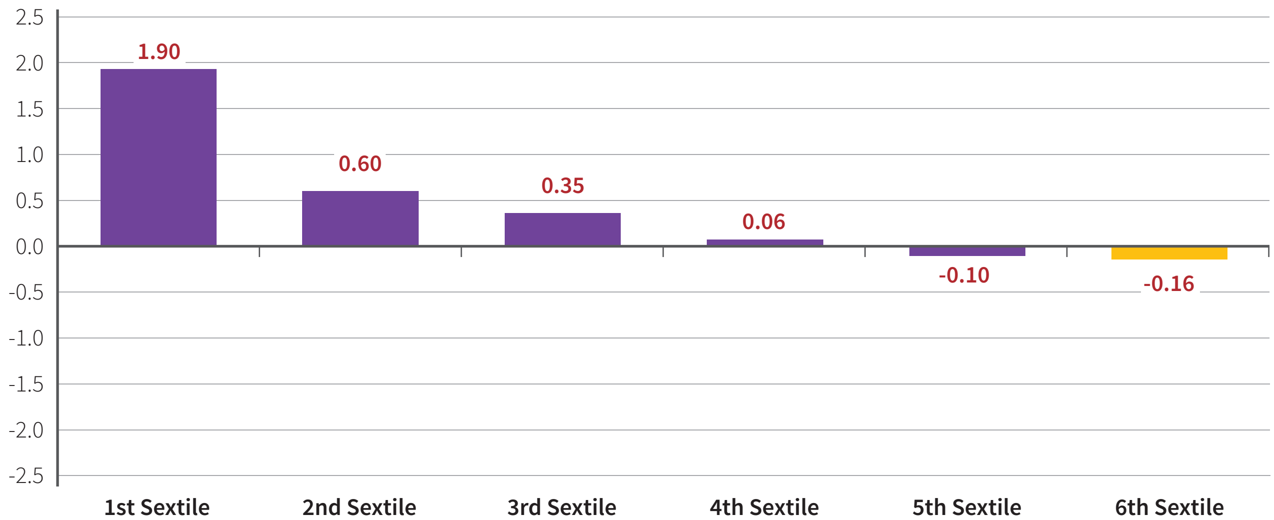

Diagram 12: Total Information Ratios across Sextiles for the EWJ

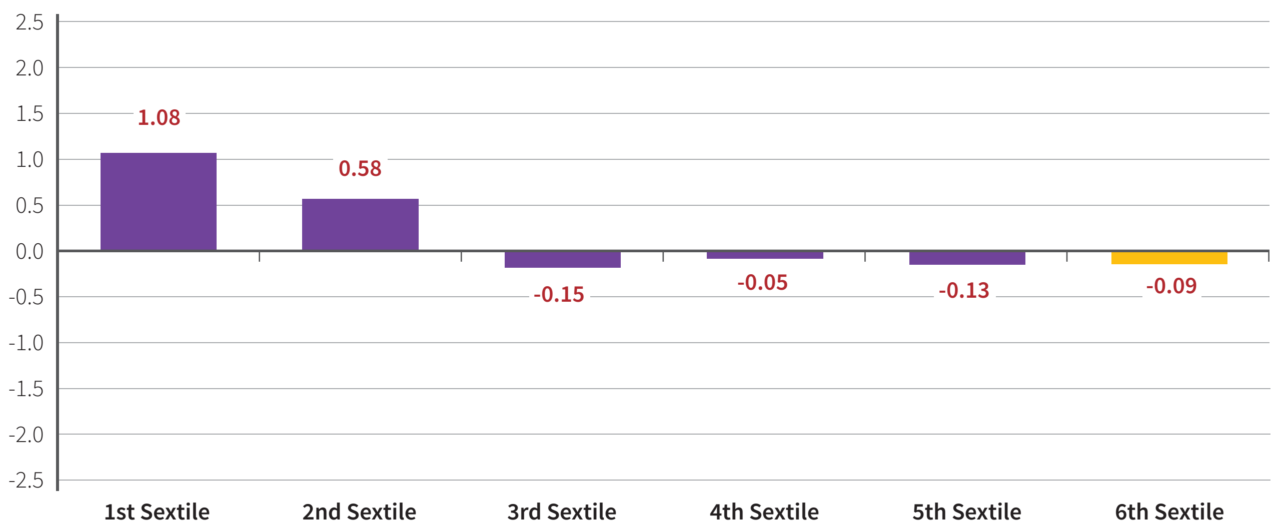

Diagram 13: Total Information Ratios across Sextiles for the EWH

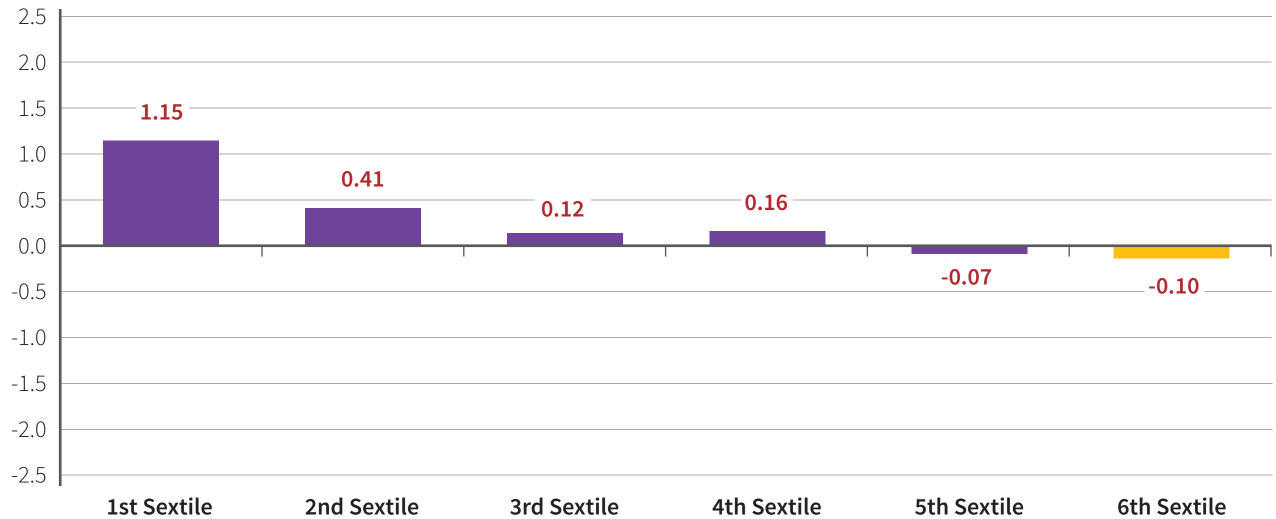

Diagram 14: Total Information Ratios across Sextiles for the EWS

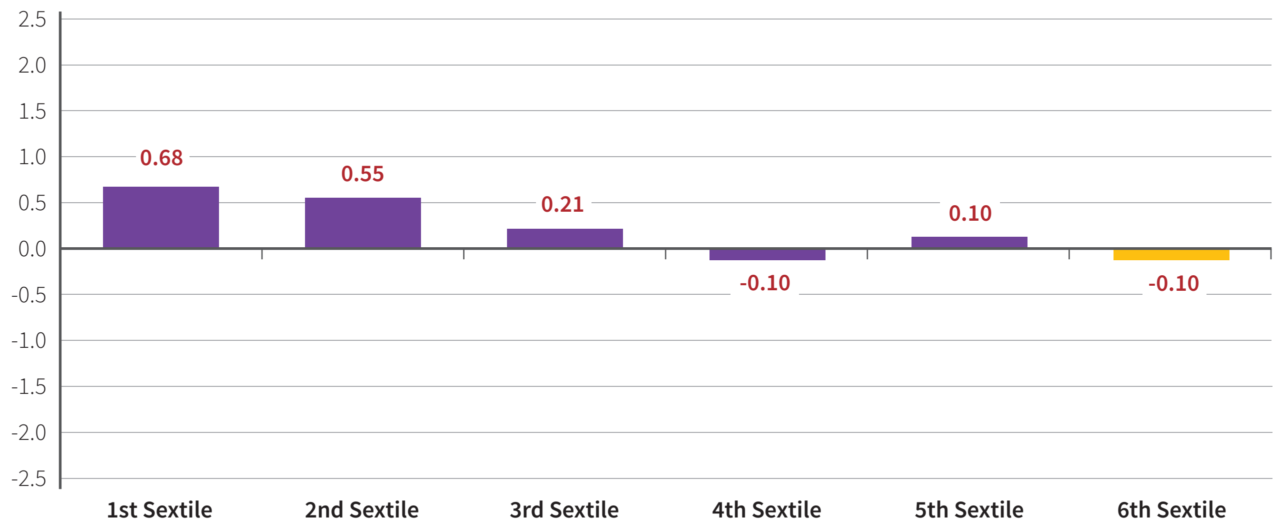

The amplitudes of the bars and the shapes of the profiles across sextiles differs somewhat but the overall pattern across all seven ETFs is very similar: Low volatility yields high compensation for the risk taken and high volatility is uniformly bad for equity returns.

And What Do We Do With These Insights?

The implication of these observations is that it is worthwhile to formulate investment strategies for equities, which, in a systematic fashion, aim to get investors out of the market during the most turbulent times.

The analysis in this note is based on hindsight and the construction of an investment strategy that would benefit from these insights is fraught with major uncertainties. In this note we took 'high volatility' to mean the upper cut-off volatility for the fifth sextile but one can only guess what the appropriate value for defining high volatility should be going forward. An equally large, or perhaps even greater, problem is that since investors will want to get out of the way of high volatilities before they materialize, they must predict short-term volatilities, with all the problems that doing so entails.

Nevertheless, any equity strategy that manages to avoid the most turbulent periods, should, when compared to a traditional buy-and-hold strategy, over a longer time period have a good chance to provide investors with improved returns with a lower level of risk.

At the very least, investors in such a risk-controlled strategy should be able to sleep better!

Disclosures and Disclaimers

CataMetrics Management

CataMetrics® Management, LLC, (CMM) is a Registered Investment Adviser which offers global, risk-managed, multi-asset class, and index-focused portfolio strategies to the institutional platform marketplace as well as to financial advisors. Please visit our website CataMetricsManagement.com for more information and to review our firm's Form ADV Part 2A.

This paper is for informational, illustrative and educational purposes only and is not intended as a recommendation of any security, or investment strategy. All investments carry a certain risk, and there is no assurance that an investment will provide positive performance over any period of time. An investor may experience loss of principal. Investment decisions should always be made based on the investor's specific financial needs and objectives, goals, time horizon, and risk tolerance. The assets and/or investment strategies described may not be suitable for all investors. Investors should consult with an investment adviser to determine the appropriate investment strategy for them.

Information obtained from third-party sources is believed to be reliable but is not guaranteed. CMM makes no representation regarding the accuracy or completeness of information or judgments provided herein. All opinions and views constitute the judgments of CMM as of the date of writing and are subject to change at any time without notice.

The Risks of ETFs

Exchange-traded funds (ETFs) are subject to risks similar to those of stocks, such as market, interest rate, foreign exchange, and liquidity risks. An investor in ETFs may bear indirect fees and expenses charged by the ETFs in addition to their direct fees and expenses, and is subject to the risk of loss of principal. ETF sponsors may suspend trading in ETFs and may not honor redemption requests. ETFs may trade at a discount or premium to their net asset value and are subject to the market fluctuations of their underlying investments. When considering investing in an ETF, you should consult your financial advisor and accountant on how investing in the fund will affect your taxes.

Before investing in an ETF, you should read both its summary prospectus and its full prospectus, which provide detailed information on the ETF's investment objective, principal investment strategies, risks, costs, and historical performance (if any). The SEC's EDGAR system, as well as Internet search engines, can help you locate a specific ETF prospectus. You can also find prospectuses on the websites of the financial firms that sponsor a particular ETF, as well as through your broker.

Past performance is no guarantee of future results of any ETF.

CataMetrics®
MANAGEMENT, LLC

55 Greens Farms Road, Suite 200-02, Westport CT, 06880 | 203-680-0330
www.CataMetricsManagement.com