Rocaton

INSIGHTS

Wag the Dog: Why Investors Should Understand and Care About Tail Risk

February 2017

203.621.1700 | **rocaton.com**

EXECUTIVE SUMMARY

As the Boy Scouts' pledge suggests, "Be Prepared." Investors should all expect to experience a significant tail event in their lifetime. Too often we hear the commentary that 1-in-100 events occur more often than we expect. Accusations abound that risk models have underestimated the probability of such events. We would suggest that it is our responsibility as practitioners to better define, understand, educate, and use risk models and forecasts to help investors prepare for future tail events.

As the investment community veers from relying solely on traditional measures such as standard deviation to gauge risk, a number of additional measures have been added to the list which serve to fortify risk analysis and investors' understanding of risk. One type of measure, which we will further explore, is well-known within and outside of the investor community and goes by a variety of names. The measure is simply an assessment of the magnitude of past or potential downside outcomes for some event at a given level of probability. We will focus on low probability (1%) outcomes and refer to the measure as the "1-in-100 outcome." Others may describe such outcomes as "first percentile outcomes" (or possibly 99th percentile if larger values are worse), tail events, or may use value-at-risk ("VaR") terminology.

While this type of measure is well-known, we believe it isn't necessarily well-understood or appropriately utilized. This can lead to imprudent dismissal of the outcomes indicated by the measure as it is, indeed, a low-probability measure on the surface. The severity of such events alone warrants a closer look. Between 1928 and 2016, the 1-in-20 downside one-year S&P 500 return outcome was -24.4%; as bad as that is, it has been almost twice the decline of -46.9% at a 1-in-100 level.[1]

Because we believe it is a key measure to incorporate in decision-making and analysis, this paper seeks to explain a so-called 1-in-100 outcome and why we think such outcomes matter in decision-making.

What does the 1-in-100 Outcome Actually Mean?

A 1-in-100 outcome represents an event for which the outcome (or a worse outcome) would be expected to occur with 1% probability in any given trial or draw. For instance, suppose we have a jar with 1,000 marbles, each of which has a different integer payoff between -$400 and $599 written on it. If we were to randomly draw a marble from the jar and were to receive a payoff equivalent to the number on the drawn marble (or pay if negative), then assuming each marble is equally likely to be drawn, the 1-in-100 outcome is a payoff of -$390.

1 Based on monthly rolling observations.

The above example is fairly simple because it represents a single random draw with simple characteristics. The nature and implications of a 1-in-100 outcome are highly dependent on the characteristics of the draws, for which there are typically three key parameters in investment return analysis[2].

1. The first parameter is the length of time that each draw represents (time period of draw).
2. The second is related to how draws are pulled, whether over discrete non-overlapping time periods or continuously or something in between (nature of draws).
3. The third is the overall time horizon of interest.

In thinking about a 1-in-100 outcome and making inferences from it, these parameters must be explicitly specified and should be understood by its users. It is very important to understand that a 1-in-100 outcome does not necessarily mean that the outcome is expected to occur once in every one hundred years.

Time Period of Draw

It is necessary to establish the relevant time period for each draw. Clearly, the probability of a given poor equity market return over a one-day period is by no means the same as the probability of that same poor return over a one-month period. Similarly, the 1-in-100 outcome over a single day is not necessarily the same as the 1-in-100 outcome over a single month. This is illustrated in **Table 1** below for observed S&P 500 returns.

Table 1: S&P 500 total returns from January 29, 1988 through December 31, 2016.

1/29/1988 to 12/31/2016	Daily Observations	Monthly Observations (Calendar Months)
% of Returns < -5%	0.2%	8.4%
1-in-100 Outcome	-3.0%	-10.8%

We can also see that it is not necessarily obvious how to infer the monthly results from the daily results. The same applies to inferring annual results from monthly results and so on. How we might do this could vary depending on the asset class and how independent we believe one period of observations is from another.

2 While we focus on 1-in-100 outcomes in this paper, the points raised are relevant to this type of measure (1-in-X, Yth percentile, etc.) in general.

Nature of Draws

Note that in **Table 1** the observations are sampled in a particular way. The daily observations rely on returns that occurred over non-overlapping one-day periods (ending with each market close) and the monthly observations are non-overlapping one-month periods that cover each calendar month.

This is especially important because in the real world, we make observations somewhat continuously and are not limited to focusing on discrete periods of time. For example, in relatively stable or normal environments, investors may review performance over trailing three-month, one-year, three-year, etc. time periods at the end of every calendar month, quarter, or year. Yet in cases where markets are "misbehaving," investors tend to watch more closely. The implication of more continuous observations is that the frequency with which we draw data increases relative to a world in which we only observe data over discrete and somewhat arbitrary periods of time (e.g., end of each calendar year). Clearly, the more frequently investors review performance relative to the frequency based on which a certain 1-in-100 estimate is calculated, the more observations or draws increase the potential for such a 1-in-100 outcome to be observed. Given that, there is a greater (or at minimum equal) probability of observing a certain pre-specified downside outcome over the course of a one-year period at least once with greater frequency of observation. This is not to say that with any specific draw the probability of the outcome necessarily changes but rather that we are increasing the number of draws over a pre-specified time period and allowing for the number of observed low-probability outcomes to increase. **Table 2** illustrates this concept with one-year S&P 500 returns.

Table 2:
S&P 500 returns from January 1928 through December 2016.

One Year Returns	Calendar Year Observations	Rolling Monthly
% of Observations (Draws)	89	1,057
# of Returns < -20%	6	69
% of Observations < -20%	6.7%	6.5%

Because the probability of an observed return of less than negative 20% over a one-year period is roughly the same with any draw, clearly, the number of times you would record such an outcome increases if you draw with greater frequency[3]. This is seen in the table above where there are 69 returns of negative 20% or worse when using rolling 12 month periods and only six such returns when using calendar year periods.

3 We will note that non-overlapping observations may not necessarily be independent from one another, and overlapping/rolling observations by definition are not independent from one another, an important observation, but outside of the scope of this paper.

Time Horizon of Interest

In addition to the time period of the draws themselves, we must also be aware of the time horizon of interest over which draws are pulled. In the above example, the time horizon of the draws was one year, but the overall time horizon explored was 89 years. Not surprisingly, increasing the overall time horizon in which draws are made increases the probability that at least one pre-specified downside outcome can occur (or, at the very least, the probability cannot decrease).

In the 1-in-100 outcome framework, to illustrate this point, let's assume that the 1-in-100 outcome for any single year return for some asset class is -10%. If we are deciding whether or not to make a 10-year investment in this asset class, then we understand that there is clearly at least a 1% chance that we will observe a -10% return or worse in at least one of the next 10 discrete one-year periods. In fact, if we assume that the returns during discrete one-year periods are independent from one another, the probability that we would observe a -10% return or worse at least once in the 10 observations is about 9.6%[4]. We also know from the discussion above that if we didn't limit ourselves to looking only at the returns from the 10 discrete non-overlapping one-year periods, the probability of observing a -10% return or worse over any one-year period within the next 10 years is indeed even higher than that.

Parameter Interaction—Drawdowns

We can extend the above issues even further if we explore drawdowns. If drawdowns (peak-to-trough returns) are relevant to us as an investor, we should understand that they are not randomly observed over a specific time interval as discussed above, but rather that the time interval over which a drawdown can occur can vary widely. Because a drawdown can occur over a variety of intervals, the likelihood that you experience a drawdown at any time over the next 10 years that is as bad as or worse than the 1-in-100 one-year outcome is indeed even higher.

If we define X as the 1-in-100 return over a calendar year:

Probability of X OR WORSE over a calendar year = 0.01
<=
Probability of X OR WORSE in any calendar year over the next 10 calendar years
<=
Probability of X OR WORSE in any one-year period over the next 10 calendar years
<=
Probability of drawdown of X OR WORSE within the next 10 calendar years

4 Calculated as 1- (1-0.01) ^10.

Figure 1:
Bringing it
All Together

Historical calendar year returns show quite a lot of volatility and downside risk. Going back almost 30 years, we have observed returns as bad as -22.1% in 2002 and -37.0% in 2008. But these observations don't quite give a full depiction of risk.

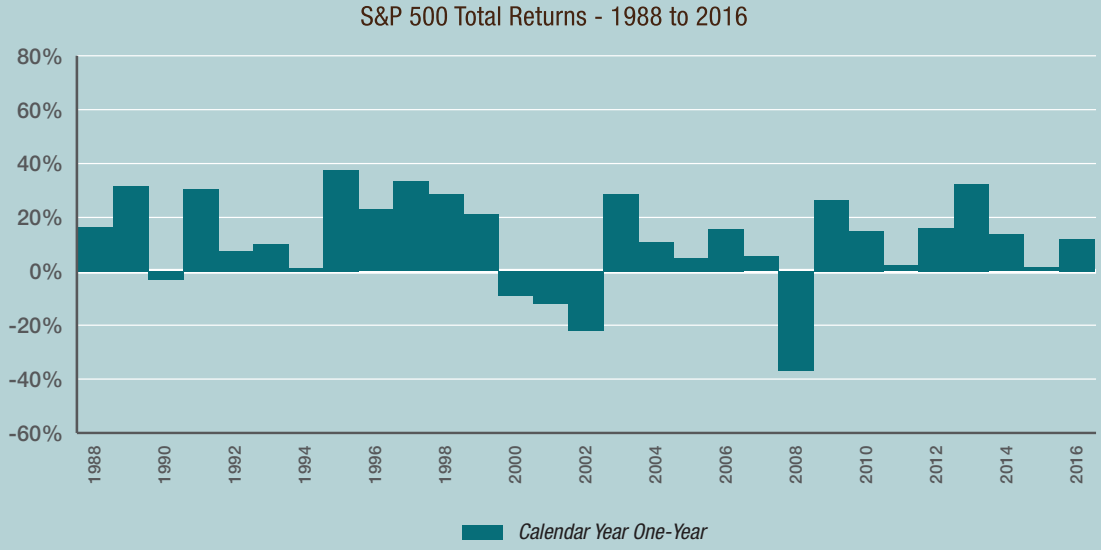

Downside returns over any one-year period (red line) will always be at least as bad as or worse than downside returns over specific discrete, non-overlapping one-year periods (teal bars), the latter of which are often the conditions under which 1-in-100 outcomes are specified, including in typical Rocaton forecasts.

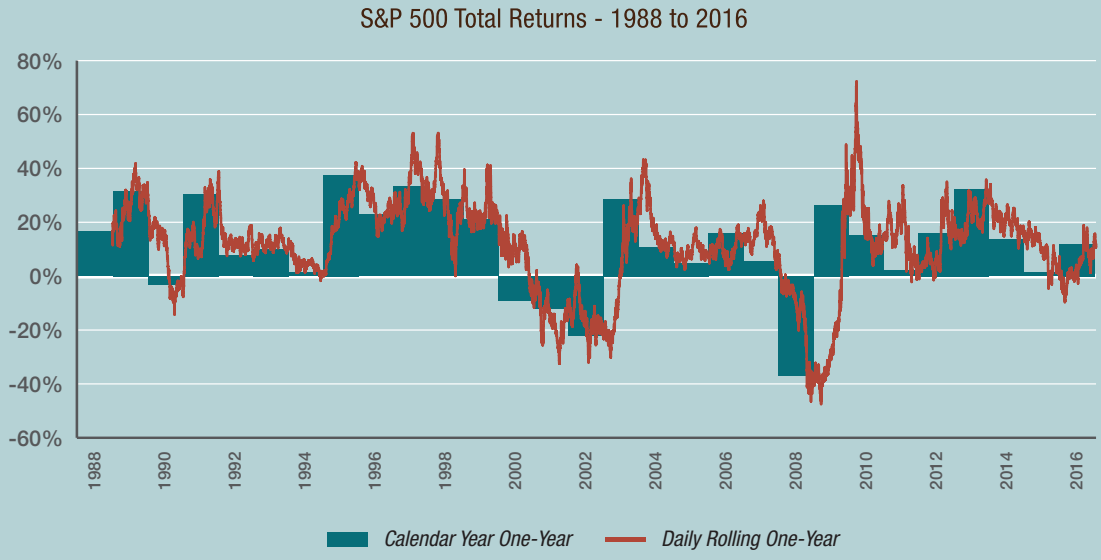

Further, downside draw-downs (green line), which occur over non-specified periods of time will always be at least as bad as or worse than downside returns over specific periods of time (red line, one-year periods).

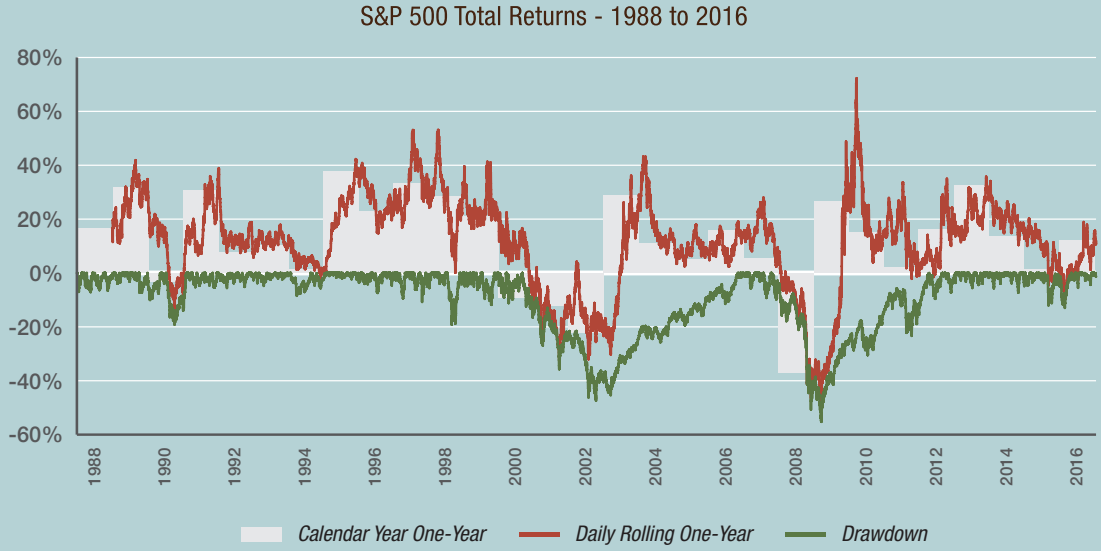

Conclusion: Bringing it All Together—Why do We Care About the 1-in-100 Outcome?

We believe that in exploring potential future outcomes, investors should take into account the 1-in-100 outcome as part of their decision-making toolkit, not only because it represents a potential outcome that should be acceptable, but also because to an extent it should be expected, given the frequency with which investors review performance and the long careers many of us expect to have in this business. This is for a variety of reasons, such as:

1. The 1-in-100 outcome used represents a shorter time horizon than an investor has; therefore, investors will have many opportunities over an extended time horizon to experience this outcome, (e.g., using a one-year 1-in-100 measure for a 10-year decision, resulting in an underestimation of the probability of observing that event over the relevant time horizon).
2. Observations can be made with different frequencies than built into the 1-in-100 calculation (e.g., we might look at one-year returns every month, not every year).
3. The 1-in-100 outcome represents a distinct and discrete period of time, but human observation occurs more continuously (especially when we are identifying bad outcomes). Drawdowns, for instance, are not measured over particular periods of time but can be a good measure of risk.

Rocaton looks forward to discussing the topics presented in this paper and continuing our conversations about risk.

Rocaton

Wag the Dog: Why Investors Should Understand and Care About Tail Risk

Disclosures

Rocaton is registered as an investment adviser with the U.S. Securities and Exchange Commission. Rocaton's Form ADV, Part 2 is available upon request. The information included in this publication has been taken from sources considered reliable. No representations or warranties are made as to the accuracy or completeness of this information and no responsibility or liability (including liability for consequential or incidental damages) is assumed for any error, omission or inaccuracy in this information. This information is subject to change over time. This publication is not intended as investment or actuarial advice. Before acting on any information contained in this material you should consider whether it is suitable for your particular circumstance and consult with your actuary. Any opinions expressed in this publication reflect our judgment at this date and are subject to change. No part of this publication may be reproduced or redistributed in any manner without the prior written permission of Rocaton Investment Advisors, LLC.

Performance Information and Return Expectations

The analysis contained in this document may include projections of long-term return and risk expectations. There is no guarantee that the projected returns or risk will be realized. The projections are based in part on historical performance of various asset classes, and past performance is no guarantee of future performance. The projections include assumptions, including those regarding risk and return. These assumptions are used for modeling purposes only and may not be realized. Because the analysis is based on assumptions and projections, there can be no warranties or guarantees.

203.621.1700 | **rocaton.com**