

MANAGING TAIL RISK

Stephen C. Sexauer and Laurence B. Siegel¹

Published in the *Retirement Income Journal*, Volume 2, Number 3 (Fall 2012)
Reprinted with permission of the authors.

EXECUTIVE SUMMARY

How should investors think about managing tail risk — the risk that very bad events will occur faster and more often than expected? Is tail risk something new, reflecting a newly interconnected world? Is it a significant factor in capital markets? If it is, what can be done to manage an investor's exposure to it? This white paper examines:

- Causes of fat tails in market outcomes
- Common misconceptions about “normal” market returns (real life is complicated; why are our models so simple?)
- Three rules for managing tail risk in investment portfolios — specifically within qualified default investment alternatives (QDIAs) for defined contribution retirement plans

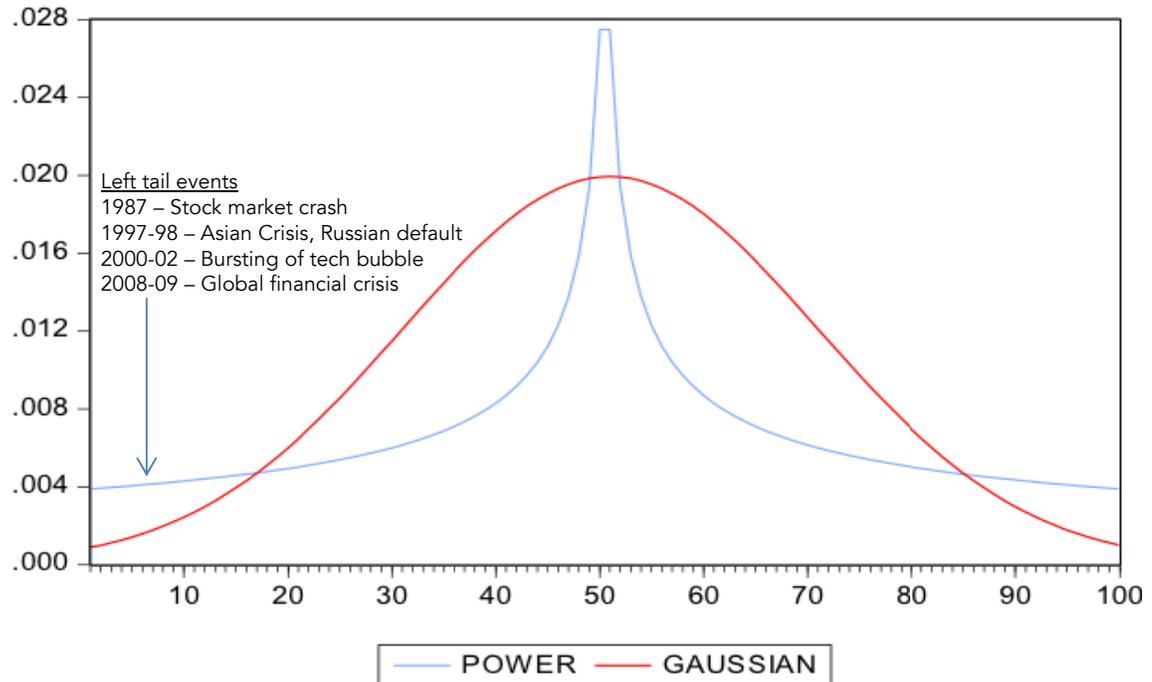
WHAT IS TAIL RISK?

Tail risk first became a hot topic after the publication in 2007 of Nassim Taleb's *The Black Swan*, which brought the obscure statistical concept of fat-tailed distributions into the mainstream.² Taleb argues that we underestimate both the likelihood and the impact of unusual events. A common way of expressing this idea is that we seem to have a 100-year flood every three or four years. This principle applies to many aspects of life (including actual floods) but is particularly relevant to understanding capital market returns, which are fat-tailed. Figure 1 compares a fat-tailed (“Power”) return distribution to a normal one (“Gaussian”). In the fat-tailed distribution, large negative returns occur several times more frequently than in the normal one — the “big one” happens more often than people believe is possible. The same is true of large positive returns, but we care more about the negative ones.

¹ Stephen C. Sexauer is chief investment officer, U.S. multi-asset, at Allianz Global Investors in New York. Laurence B. Siegel is the Gary P. Brinson director of research at the Research Foundation of CFA Institute, Charlottesville, VA, and senior advisor to Ounavarra Capital LLC in New York. The authors may be reached at Stephen.Sexauer@allianzgi.com and lbsiegel@uchicago.edu.

² An earlier Taleb book, *Fooled by Randomness* (2004), struck many of the same notes but did not have the same focus on unusually large downside events. It is well worth reading.

Figure 1
Fat-Tailed ("Power) and Normal ("Gaussian") Distributions Compared



Taleb's timing was either lucky or extremely good in that the capital markets responded the very next year (2008) with a black swan-type event. The S&P 500 fell by 37% for the calendar year and 57% on a daily peak-to-trough basis.³ International equities fell even more. Real estate equity, subprime mortgages, corporate bonds, treasury inflation protected securities (TIPS), and crude oil fell by amounts that were well outside the experience of the last few decades. Asset mixes that were diversified in normal times ceased to be, and correlations across asset classes moved toward 1.0 — all risk assets went down, and went down a lot. While the collapse of the capital markets did not reach Great Depression levels, the magnitude and speed of the downward moves were reminiscent of that depression and provoked widespread concern about a repeat performance.

Naturally, a wide cross-section of investors became interested in avoiding, hedging, or otherwise managing tail risk.

The crash of 2008 was hardly the only example of tail risk in investors' experience. Others include the 1997 East Asian crisis and "contagion" that began in Thailand and rapidly spread over the East Asian Tiger economies; the Russian debt default and the associated halt in ruble trading; the attacks of September 11, 2001; the 2001 default in Argentina (the largest in history prior to Greece), the bursting of the technology bubble over 2000-2002; and the current European sovereign debt crisis. And that's just the last 15 years!

³ The peak (S&P at 1565.15) was on October 9, 2007, and the trough (S&P at 676.53) was on March 9, 2009. These numbers are daily closing prices.

LIFE IS COMPLICATED. WHY ARE OUR MATHEMATICAL MODELS SIMPLE?

The late French mathematician Benoit Mandelbrot pointed out in a 1963 paper that “real-life economic systems ... are ... dominated by ... the extreme cases. Specifically the outer 5 percent are...as important as the [other 95%] of the data.”⁴ Mandelbrot is best known for his beautiful computer drawings of fractals — geometric objects that are “self-similar,” looking the same at different scales like the outline of a snowflake. The fractal distributions used to produce these pictures have infinite variance, a very uncomfortable property for financial analysts who are used to dealing with a finite, measurable variance.

Mandelbrot concluded that markets follow something like a fractal distribution and can have infinite variance. Like snowflake outlines, financial market returns are self-similar; a stock index chart looks pretty much the same at yearly, daily and minute-by-minute scales. This does not mean that market returns have infinite variance but it sure is a hint that they’re not normally distributed. At the very least, market return distributions are fat-tailed (leptokurtic), with bad events happening more often than predicted by the normal or Gaussian distribution. Big changes in markets happen fast, and far more often than “normal.”

Yet common practice in finance is to assume that market returns are normally or log-normally distributed (the classic or Gaussian bell curve), with the effect that fat tails — which really do exist — are assumed away. The principal reason is that, under a normal distribution assumption, the math is easy. You can calculate volatility, alpha, beta, the Sharpe ratio, and other important measures. All of science involves making simplifying assumptions so that workable tools and heuristics can be used to better understand nature, which is complex and unpredictable. Normal distributions make for workable tools and heuristics.

WHAT CAUSES FAT-TAILED OUTCOMES?

Markets reflect economic life, with market prices being determined by (1) expectations for growth in the real economy and (2) the current set of discount rates, one for each type of asset. Of these two factors, the real economy is by far the more important in the long run.

Economic growth can come from only two sources: an increase in productivity per worker (innovation) and an increase in the number of people working; we can ignore the latter because it is mostly predictable. Innovation, the source of productivity gain, is lumpy: A google is rare and big, both as a firm (Google) and as a number (10^{100}). The Internet flattens the earth once, and no one can unflatten it after that. Disruptive technologies — the railroad, the automobile, the electrical grid, the medical revolution — are just that; they disrupt the existing order, making some people unexpectedly rich, putting others out of business and moving markets profoundly.

⁴ Paraphrase by John Matson, “Benoit Mandelbrot and the Wildness of Financial Markets,” *Scientific American* news blog, March 13, 2009. Mandelbrot’s paper, “The Variation of Certain Speculative Prices,” appeared in *The Journal of Business*, Vol. 36, No. 4 (October 1963), pp. 394-419.

On the negative side, transient events, such as natural disasters, can have a huge economic impact: The effects of Hurricane Katrina, the Indian Ocean tsunami, and the Japanese earthquake and tsunami are still being felt. Political and policy changes—human decisions—can overwhelm all other factors. When the rules and incentives change for better or worse, they can do so quickly and violently—as we’ve seen in Russia, Venezuela, Iran, and Egypt among many other places. We’re now facing the potential breakup of the euro, an event with consequences that are unforeseen by definition because the euro was designed to last forever. (We beg our readers’ pardon for chuckling at that concept.)

Markets, thus, reflect people and their unpredictability. Fundamental values of securities and markets can change all of a sudden. Perceived values can change just as quickly. Markets are not “normal.” Markets have tail risk.

MANAGING TAIL RISK IN PORTFOLIOS

Managing tail risk is like managing any other market-related risk. It requires a qualitative and quantitative understanding of (1) the sources of the risk, which we’ve just covered, and (2) the effects of the risk, how it affects market prices and returns. We can express the principles of managing tail risk in three rules:

- **Rule #1: Be aware. Tail risk exists.** That is, the normal distribution is not up to the task of predicting the range and frequency of returns. Specific tail risks (negative events) are exceedingly hard to predict, and therefore to hedge and protect against.
- **Rule #2: To manage tail risk, you have to enumerate possible outcomes.** The outcomes can be specific events (euro collapse, war in Iraq, energy price spike, etc.). Estimate, or guesstimate, their effects. Develop conceptual statistical models of them. (Are the fat tails part of a leptokurtic distribution? A bimodal one? A negatively skewed one?) Awareness plus judgment matter most. Don’t allow your investment strategy to depend on the world being smooth and Gaussian—it won’t be.
- **Rule #3: Since you cannot avoid tail risks—they are part of the texture of the world—you have to choose specific strategies to manage the risks.** There are three basic approaches:
 - a. *Low-beta strategy.* Take less of the risk.
 - b. *Passive hedge.* Pay to constantly hedge the large bad events. (The simplest example of this strategy is to have a permanent long-volatility position. This is very expensive.)
 - c. *Active hedge.* Put the hedge on when you think that risk in general is about to rise, or when you fear a specific bad outcome. Take the hedge off at other times.

When losses cannot be tolerated, they should be passively protected against with a constantly reviewed low-risk or low-beta asset allocation. This protection comes at a cost, which is a lower expected return and thus a higher required savings rate. When losses can be tolerated (tail-risk hedges do not work all the time) and the manager has specialized skills in identifying future tail-risk events and forecasting their changing probabilities over time, tail-risk hedges can be purchased that should decrease the severity of downside events and increase returns.⁵ The total cost includes the direct cost of the hedge, the cost of the specialized skills, and the cost of occasional failures. There is also some loss of risk transparency because many tail-risk hedges are indirect hedges, typically using derivatives in the futures and options markets. (The hedges themselves are deemed “speculative” from a regulatory standpoint.)

MANAGING TAIL RISK IN QDIA PORTFOLIOS

A reasonable expectation for qualified default investment alternatives (QDIA) portfolios is that they be transparent and reliable, and that the process used to construct them be repeatable over a period of many years (up to 40!) and across a wide range of market cycles. The most transparent, reliable and repeatable way to manage tail risk is passively, using a low-beta structure when large losses cannot be tolerated. An example of a time when large losses cannot be tolerated is the “sequencing risk” period in target-date funds (TDFs) the decade or so just before and after retirement.

The most important QDIA target-date fund decisions are (1) designing the participant choice architecture, (2) evaluating and picking a QDIA target-date fund risk schedule (the glide path) and a TDF manager, and (3) monitoring the results of the entire program. These decisions and ongoing tasks are performed by the plan sponsor with input from the sponsor’s advisor. Within the plan sponsor’s decision of a TDF manager will be the choice of how tail risk is viewed and managed. The important points here are (1) that it is the plan sponsor who is controlling the process and making the decisions; and (2) issues of tail-risk hedging, while significant, are dominated by the higher-order decisions of plan design and the choice of the appropriate risk schedule.

The specific choice of active or passive tail-risk management will be customized to the needs of the plan participants, the ability of the plan’s trustees to evaluate and monitor the portfolio managers, and the portfolio management capabilities available in the market at that point in time.

Explicitly managing tail risks — either through a passive hedge or active management — can bring benefits when hedges are successful. These expected benefits need to be balanced against (1) the reality that the hedging cannot always be successful — prudence and experience tell us to expect some failure; (2) there will be some loss of transparency when futures, options, and other hedges are added; and (3) successful active tail-risk hedging requires

⁵ There has been much work in this area. For example, PIMCO has pioneered the use of tail-risk hedging, and has written on how to apply these tail-risk strategies. See Bhansali, Vineer, “Tail Risk Management,” *The Journal of Portfolio Management*, Summer 2008.

sophisticated and dedicated professionals with expertise in markets, futures and options, and mathematics.

Tail risk is important, but it is only one of many risks with which investors and their managers should be concerned. A responsible approach to TDF selection involves awareness by the sponsor and its advisors of the range of approaches in dealing with tail risk in a QDIA program, and a commitment to do so as part of the full spectrum of risks faced by the plan sponsor and individual participant.