LOOKING INTO THE FUTURE CASTS SHADOWS

A PERSPECTIVE ON PORTFOLIO THEORY, FAT TAILS, AND RISK MANAGEMENT

John R. Minahan, CFA[1]
Senior Investment Strategist

Introduction

2008 reminded us that human beings have limited ability to peer into the future. As investment professionals we have attempted to enhance this limited ability by using various models and theories as lenses through which to form views of future risks and returns, and to select investment strategies in light of those views. Yet many of us were caught off guard by 2008. Does this mean that our lenses were cloudy, that our models were flawed? Absolutely. Indeed, I believe this is how it must be. No matter how good our models are, they will never be able to change the fact that the future is unknown. At best models are like flashlights one can shine on a narrow swath of the future, a swath that casts shadows.

One way of shining a light on the future is to assume that the aspects of the future you care about (return, usually, for investors) can be represented by a probability distribution, that is, by a list of all possible values a variable can take on and an assignment of relative likelihood to each possible outcome. Investment professionals often do this by assuming investment returns are normally distributed and have a well-behaved correlation structure. In recent years, however, we have been repeatedly reminded – both by events and by authors[2] shining lights from different angles – that investment returns have neither of these properties, and in fact display what have come to be called "fat tails" and "black swans".

This paper is an introduction to portfolio risk management in the presence of fat tails and black swans. We have four primary objectives:

- To provide the reader an intuitive understanding of the concepts of fat tails and black swans, and why these concepts present a challenge to traditional portfolio theory.
- To paint a balanced portrait of portfolio theory and to evaluate various popular criticisms of the theory in light of that portrait.
- To define risk measurement from first principles – principles general enough to encompass normal distributions, non-normal distributions, and no distributions at all.
- To suggest some principles of investment risk management in the presence of fat tails and black swans

Fat Tails, Black Swans, and the "Failure" of Portfolio Theory

Investment returns typically display extreme outcomes more often than would a normal distribution with the same mean and variance. This phenomenon is alternately called "fat tails" or "excess kurtosis" depending on the context and who is talking[3]. See Exhibit 1 on the following page for a picture of distributions with various amounts of kurtosis.

A related but much broader issue is the truism

[1] I thank Adam Apt, Elizabeth Chertavian, Mark Cintolo, Dan diBartolomeo, Barclay Douglas, Chris Levell, Erik Knutzen, Larry Pohlman, and Susan Weiner for helpful comments and other assistance.
[2] See Bookstaber (2007) or Taleb (2007), for example.
[3] Both terms derive from statistics. A fat-tailed distribution is one which exhibits extreme outcomes more often than a normal distribution with the same mean and variance. Kurtosis is a specific quantitative measure of how fat-tailed a distribution is. The normal distribution has a kurtosis of 3; any amount more than this is called "excess kurtosis". A detailed description of kurtosis is in contained in Exhibits 1 and 2.

that *all* models of the future will emphasize some things and de-emphasize others, and will therefore cast shadows which can create blind spots for users of the model. Sometimes problems can grow in these blind spots, until they become big enough to shift attention to themselves, at which time they are experienced as "black swans", or unprecedented, hard-to-imagine, extreme events[4,5]. Black swans and fat tails are closely related: a black swan can be thought of as a fat tail that hasn't happened yet. Please see Exhibit 2 (in appendix) for a detailed example which illustrates the relationships among fat tails, black swans, and kurtosis.

Exhibit 1 - Family of Distributions with Different Kurtosis

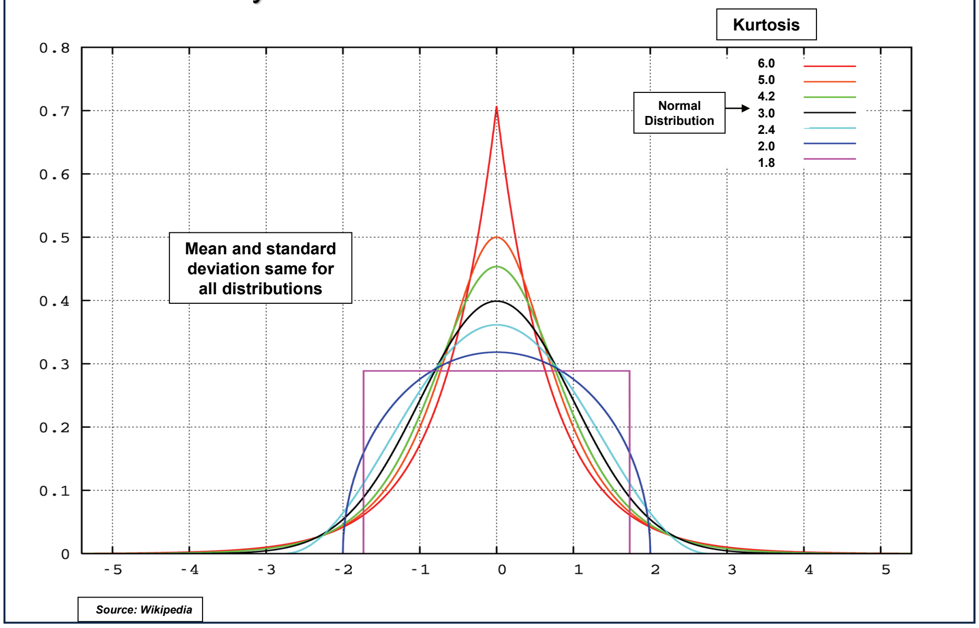

Greater awareness of fat tails and black swans – intensified by the market crisis[6] – has brought into question the relevance of normal distributions and the portfolio and capital market theories built upon them. Investors are reexamining the processes by which they made decisions that resulted in losses far beyond expectations. The image of mean-variance portfolio theory as a failed risk management framework is a common theme in these accounts. New approaches to risk management are being called for, approaches which explicitly account for fat tails, black swans, and "downside risk".

II. Post-Modern Portfolio Theory

I have had many discussions in the months since the financial crisis began – in client meetings, over lunch, in the hallway, at conferences, at parties – about the future of mean-variance portfolio theory. In these discussions, it is not unusual for someone to wonder out loud something like, "Perhaps it is time to finally put the nails in the coffin of modern portfolio theory".

I've always been a little confused by "modern" portfolio theory – the portfolio theory I learned in school was developed in the 1950s and might better be called "classical" portfolio theory, although just "portfolio theory" is sufficient, in my view. Nonetheless, when I hear people saying it is time to move beyond it, they almost always call it modern portfolio theory or just "MPT". To such people I say: "I find portfolio theory extremely useful for framing asset allocation decisions. Portfolio theory does indeed have shortcomings, both real and imagined, but these are well known and can either be compensated for or avoided."

[4] Merton (2008) put it something like this: the blind spots are not priced, and therefore provide an opportunity to arbitrage the difference between models and reality for someone who sees both; bringing financial engineering to bear on such arbitrage causes a bubble – a mispricing which feeds on itself – at least early on when it appears the value-opportunity is real. If the bubble grows unchecked it will eventually blow up the market and what had been an unappreciated difference between the model and reality becomes a (now-fully-recognized) black swan. This idea is similar to Bookstaber's (2007) idea that tight coupling combined with complexity will inevitably lead to an accident.

[5] The term "black swan" derives from the fact that, for hundreds of years, all swans observed by Europeans were white. Philosophers would use the idea that a black swan might exist to illustrate the point that, just because you haven't seen it doesn't mean it doesn't exist. And indeed, blacks swans did exist, and were eventually discovered by Europeans, in Australia.

[6] An interesting question is whether 2008 was a fat tail or a black swan. The stock market return was certainly a fat tail, not a black swan, since returns in the vicinity of -40% have been experienced before. Credit market returns and behavior may have a greater claim to being a black swan, but even this is not shocking

I find it fruitful to think of portfolio theory as a framework which can be deconstructed into its component parts, and then reassembled using the elements one considers valid in a given context, incorporating elements of other approaches which complement the weaknesses of portfolio theory. No theory is perfect. It is useful to get a couple of different perspectives on important decisions.

Exhibit 3 (in appendix) illustrates how I deconstruct portfolio theory. I lay out different aspects of portfolio theory so that individual components can be assessed on their own merits. I don't attempt to be comprehensive. The characteristics I chose to highlight are those I find myself debating with people. I hope that this grid will allow for a more focused discussion.

THE INVESTMENT PROFESSION HAS CREATED ENORMOUS CONFUSION AROUND THE MISTAKEN IDEA THAT STANDARD DEVIATION AND RISK ARE SYNONYMOUS.

I do believe there are things we need to unlearn about portfolio theory. We need to stop thinking of standard deviation as a measure of risk[7]. We need to be clearer that "expected" only has a 50/50 chance of being achieved[8]. We need to be more explicit that portfolio theory is forward looking and not dependent on historical data, though it is informed by it. And in circumstances where fat tails or illiquidity may be present, we need ways of analyzing those risks. Most importantly, we need to recognize that portfolio theory is *just a model*. Like any model, it captures some elements of reality and misses others.

Once we unlearn some bad habits regarding portfolio theory, I think we will find that the theory is quite useful, albeit incomplete.

III. Risk Measurement

Risk is the possibility of something undesirable happening[9]. In an investment context this usually means the possibility of poor returns.

Risk is not inherently measurable. Only by making assumptions about (1) an investor's ability to forecast probability distributions and (2) how the investor weights the displeasure experienced from various bad outcomes, does risk become measurable – *to the specific investor in question*. Unless one assumes all investors are the same, there is no basis for defining a risk measure as characteristic of an investment rather than as a characteristic of a specific investor's evaluation of that investment.

Standard deviation is a useful summary statistic of a probability distribution – it gives you a quick sense of how spread out the distribution is. And, if you also know the mean and assume the distribution is normal, you can calculate the probability of any possible future outcome.

Standard deviation is not risk. Unfortunately, the investment profession has fallen into some poor linguistic habits which have created enormous confusion around the mistaken idea that standard deviation and risk are synonymous. Some of this confusion has contributed to what seems to be a popular revolt against portfolio theory. Portfolio theory is often criticized – mocked even – because it allegedly treats upside and downside deviations from expectations symmetrically. Yet this is a problem not with portfolio theory per se but with the idea that standard deviation is risk, an idea which is not central to portfolio theory (see footnote 8 for more discussion of this point).

[7] In the original paper on portfolio theory, Markowitz (1952) develops the entire theory without using the word "risk". Then, in his concluding comments, he says "The concepts 'yield' and 'risk' appear frequently in financial writings. Usually if the term 'yield' were replaced by 'expected yield' or 'expected return', and 'risk' by 'variance of return', little change of apparent meaning would result." This throwaway comment was a tactical disaster. The idea that "portfolio theory equates variance with risk" came to be widely accepted as a cornerstone of portfolio theory, which in turn led to excess skepticism for the theory by practitioners who know that the idea that risk is variance is silly. Yet, nothing in Markowitz's theory requires that risk be equated with variance. His unfortunate comment to that effect and the subsequent development of a finance jargon which conflates risk and standard deviation seems to have led many people to reject portfolio theory for the wrong reason.

[8] Confusing the statistical and English meanings of "expected" can lead to excessively risky portfolios if the investor reaches for more and more "expected return" to meet an arbitrary target, and ignores the extent to which the left tail expands as one reaches.

[9] Knight (1928) draws a distinction between situations in which a probability distribution is known – he calls this "risk" – and situations where probabilities are not known – he calls this "uncertainty." Whether or not one knows a probability distribution is a very useful distinction, but Knight's labeling of these situations as "risk" and "uncertainty" does not correspond to everyday uses of these terms and is less helpful.

Section III.1, below, describes the relationship between standard deviation and risk when returns are normally distributed in an attempt to disentangle the "standard deviation is risk" misunderstanding. We discuss risk measurement with fat tails and black swans in sections III.2 and III.3, respectively, and address the distinction between idiosyncratic and systemic fat-tails in section III.4.

1. Risk measurement in a normal context

Assuming that returns are normally distributed does not force one to define risk in any particular way. All of the various "downside risk" measures[10] – probability of a loss, probability of a disaster (however defined), probability of failing to meet an objective, downside deviation, and all other "lower partial moments[11]" – can be calculated using a normal distribution.

Standard deviation and risk are closely related, as is illustrated in Exhibit 4:

- If one defines risk as the probability of a loss (or any lower partial moment), then risk is in part determined by standard deviation. (Exhibit 4a)
- If one holds the mean constant, then an increase in standard deviation always implies an increase in the probability of loss. (Exhibit 4b.) Similar illustrations could show that all lower partial moments increase if variance is increased while holding the mean constant[12].
- If one holds the variance constant and increases the mean, one can see why standard deviation is not a risk measure: two distributions can have the same standard deviation and different probabilities of a loss. (Exhibit 4c)

Exhibit 4a - The Relationship Between Risk and Standard Deviation

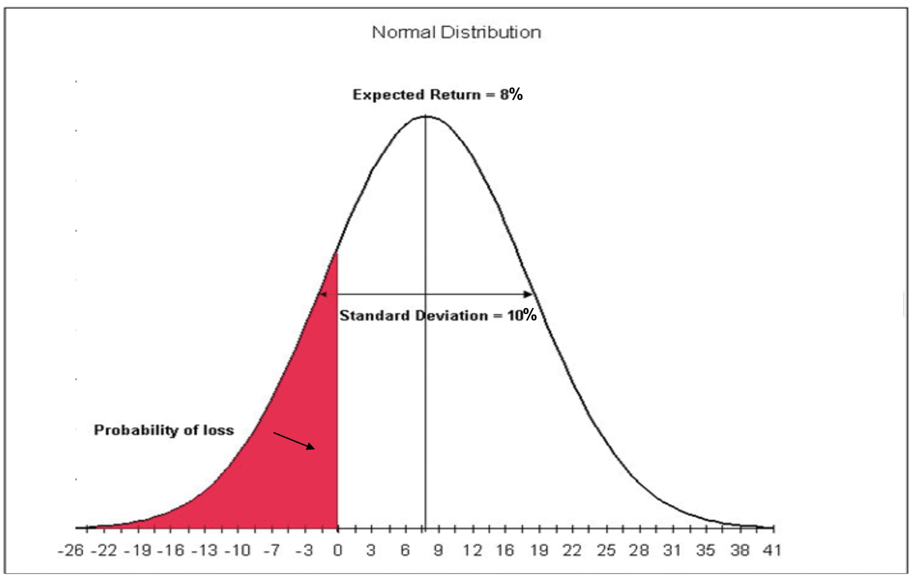

2. Risk measurement with fat tails

If one is dealing with a known fat-tailed distribution, then risk measurement is straightforward. One can simply observe the distribution and decide what one does and does not like about it. Or, if a specific measure is desired, any risk measure of interest can be calculated if the distribution is known.

In real life one almost never deals with a known distribution. At best we have a historical sample and some conceptual models for affirming or modifying the historical sample. Nonetheless, it can be useful to treat the historical distribution as if it is possibly representative of the future distribution. One must be mindful, however, that statistics calculated with historical data are merely suggestive of future possibilities.

As an example of how one can use fat-tailed historical samples to inform future expectations we explore two questions about US stocks:

[10] "Downside risk" is redundant. Risk is by definition a downside phenomenon. That we feel the need to clarify that we mean "downside" risk illustrates how confused our language about risk has become.
[11] A "lower partial moment" is a weighted average of the events below some threshold. Both the threshold and the weighting scheme are parameters which can be chosen by the risk measurer. All "downside risk measures" are lower partial moments of one type or another.
[12] Rothschild and Stiglitz (1970) show that a mean-preserving increase in standard deviation is not in general the same thing as an increase in risk. It is, however, for a normal distribution.

Exhibit 4b - The Relationship Between Risk and Standard Deviation

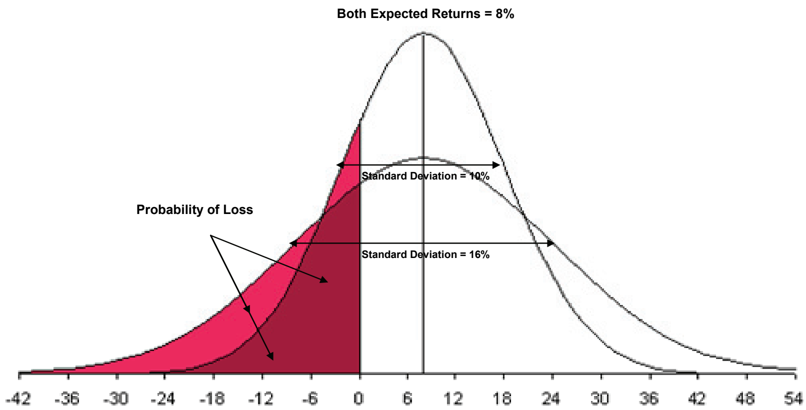

i) What is a plausible worst-case scenario for future returns on US stocks?

Suppose you wanted to estimate the worst drawdown of the S&P 500 on a forward-looking basis. If you looked at the history of the S&P 500 (See Exhibit 5), you would find that once in the last 80 years there was a fall close to 90%, and that four times in the last 35 years there has been a fall close to 40%. It would certainly be reasonable to assume on the basis of this historical data that there is a remote possibility of a 90% fall in the S&P 500 looking forward, and a pretty good possibility that we will experience another 40% decline sometime in the next decade or so. Which of these possible outcomes you treat as your worst-case scenario for planning purposes is a matter of judgment.

It really isn't necessary to translate the distribution into summary statistics since the distribution itself has the information we want. Indeed the summary statistics lose information relative to the distribution itself. For example, if you were to estimate the likelihood of a down 40% and down 90% event by calculating the mean and standard deviation of historical returns and assuming they are the mean and standard deviation of a normal distribution, you might conclude that losses of 40% occur only once every 10,000 years, and that 90% losses never occur. Fortunately, there is no reason to make this mistake. If you know the historical data, you can incorporate them into your decision process directly without having to calculate summary statistics.

Do fat tails decline with time?

The idea that "return distributions become more normal-looking the longer the measurement horizon" is suggested by the central limit theorem and has been demonstrated for US stocks by Jeremy Grantham (2009). Specifically, Grantham showed that daily stock returns have very fat tails, annual stock returns have moderate fat tails, and 30-year returns do not have fat tails. (Exhibit 6)

If fat tails do decline with time horizon, then long-term investors can be less concerned with them as long as they don't become liquidity

Exhibit 4c - Why Standard Deviation Is Not Risk

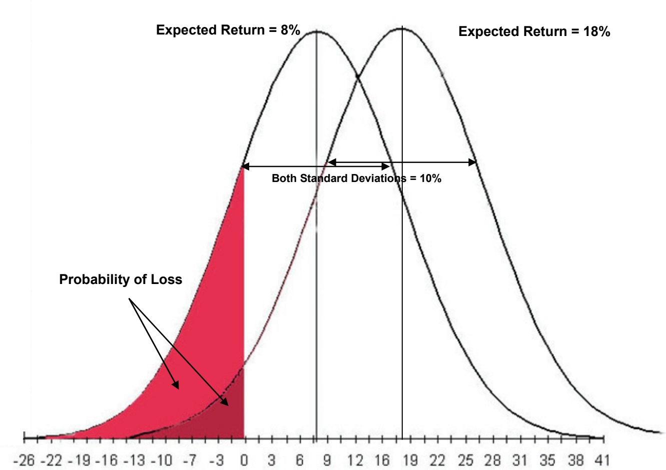

Exhibit 5 – Return History of the S&P 500

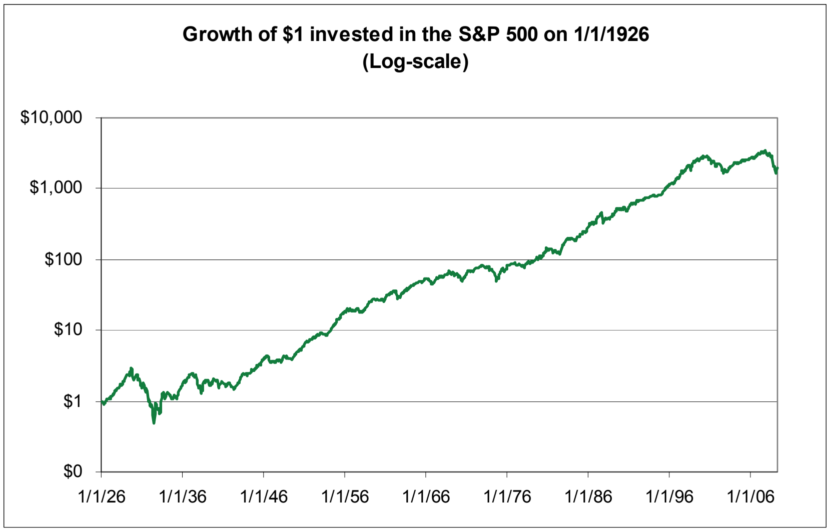

Source: NEPC Research

constrained during the fat-tailed event itself[13]. This is not to say long-term investors are not hurt by extreme left-tailed events, but that over the long-term, short-term fat-tailed events reflect themselves in the mean and variance of the long-horizon distribution, rather than in the kurtosis of the long-horizon distribution.

On the other hand, if a long-term investor becomes distressed during a fat-tailed event and is forced to trade in an illiquid market, that investor's time horizon has effectively become short, and the investor does not benefit from the lack of fat tails in the long-term distribution. The possibility of this scenario elevates forward-looking liquidity analysis to a central location in the investor's investment strategy selection process, especially for investors in illiquid assets and/or investors who use leverage.

3. Risk measurement with black swans

Black swans are unprecedented, hard-to-imagine, extreme events. All three of these characteristics are in the eye of the beholder. For example, we may all agree that the market environment of late 2008 was extreme and not imagined by most investors, but people have differing views as to whether it was unprecedented.

Black swans can be thought of as arising out of "model risk", that is, the divergence between one's model (mental or mathematical) and reality. Because black swans are invisible until they are experienced, there is not a precise way of measuring them. That said, as Bookstaber (2007) has argued, black swans are more likely to manifest themselves in complex systems.

The central risk management issue with black swans is not one of measurement, but awareness. Once one becomes aware of them, they are no longer black swans.

4. Systemic fat tails are more of a concern than idiosyncratic ones

Investors with a reasonably diversified portfolio will be primarily concerned with systemic fat

Exhibit 6a - The Wildness Lies in Wait

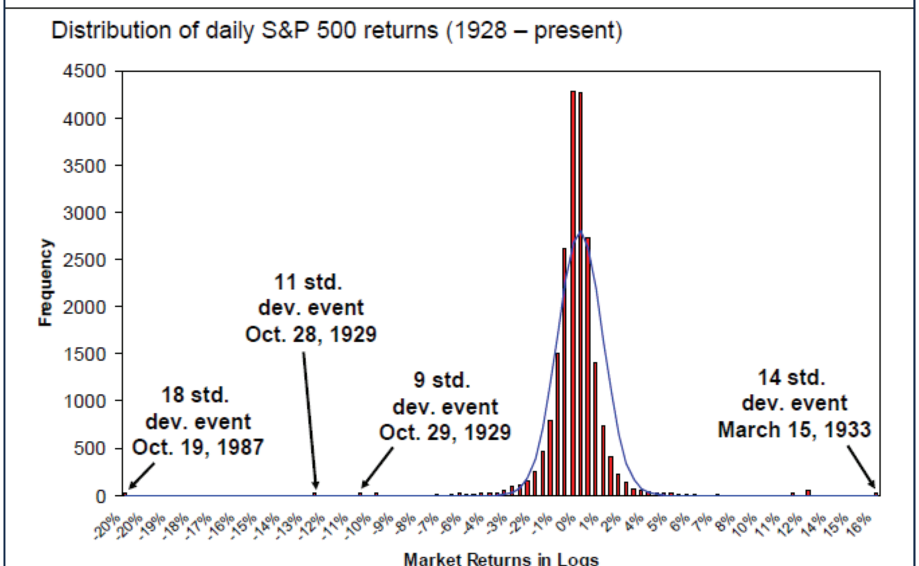

GMO Source: Standard & Poor's, Global Financial Data 0

[13] I first heard this articulated by Dan diBartolomeo of Northfield Information Services.

Exhibit 6b - The Wildness Lies in Wait: Overresponse at 1 Year

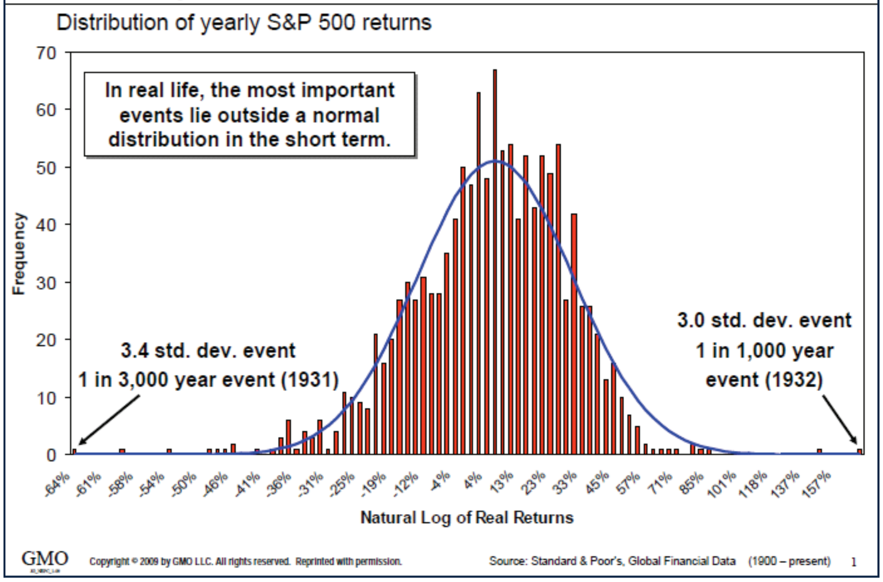

tails[14], since idiosyncratic fat tails do not aggregate to the portfolio fat tails unless the portfolio is concentrated. One way of characterizing systemic fat tails is through the phrase “all correlations go to one in a crisis”.

All correlations going to one is a short-term phenomenon. As has been demonstrated by Inker (2009), high correlations measured with monthly data do not imply that investments are tightly linked on a long-term horizon. Inker uses the example of emerging markets versus US stocks. Over the last ten years emerging markets earned 3 times the cumulative returns of US stocks, and have also displayed high short-term correlations with US stocks. Suppose you knew in advance one of these markets would earn cumulative returns three times the other, but you didn’t know which – the two markets would be prime diversifiers for each other despite high short-term correlations.

So the issue once again comes down to a question of liquidity. “All correlations going to one in a crisis” only creates a fat-tail issue for investors forced to liquidate in the crisis.

IV. Risk Management with Fat Tails and Black Swans

Risk management is the process of identifying what can go wrong with contemplated courses of action and taking steps to keep the likelihood and magnitude of unacceptable outcomes within tolerable limits. Risk management is first about awareness, second about thinking through the potential consequences of one’s actions, third about avoiding courses of action that seem too risky, and lastly about measurement.

A useful starting point for risk management is to ask the questions:

1. <u>What is the least risky course of action?</u> Attempting to answer this question forces clarification of the outcomes one wishes to avoid, and also provides a useful benchmark against which to evaluate the risk and prospective return of other options[15].

Exhibit 6c - But, at 30-Year Horizons the 1-Year Overresponse Looks Silly

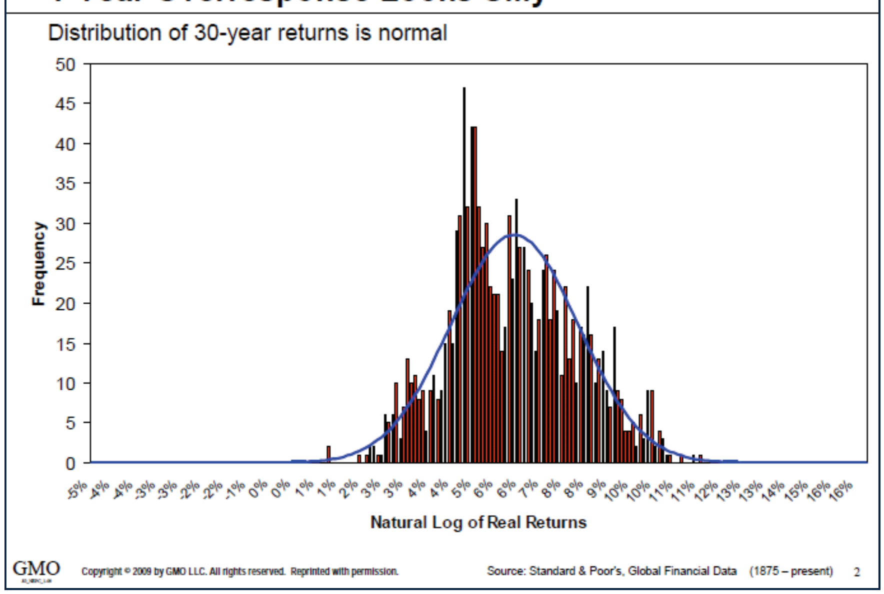

[14] Bhansali (2007) makes this point.
[15] The least risky option could be cash, a liability-hedged portfolio, or anything which corresponds to the investor’s way of thinking about risk minimization.

2. What risks remain even when one chooses the least risky option? It is important to recognize that some risks cannot be eliminated.

3. What is the likely return of the least risky option? It should be pretty easy to estimate the return of the least risky option. After all, if return is hard to estimate, then you probably aren't dealing with the least risky scenario.

4. Under what circumstances would you be willing to risk performance lower than that of the least risk option in an attempt to achieve performance higher than it? How much are you willing to put at risk? How high do prospective returns need to be to make it worth placing your assets at risk?

ALL CORRELATIONS GOING TO ONE IN A CRISIS IS ONLY AN ISSUE FOR INVESTORS FORCED TO LIQUIDATE DURING THE CRISIS.

It is useful to classify investment strategies by the extent to which they depart from the least risky alternative in pursuit of higher returns. As one moves from lower risk to higher risk options, not only does the risk increase, but one's ability to describe that risk with a probability distribution decreases. This increases the important of qualitative risk management as one moves to more risky investment options.

Although there is an infinite number of ways investment strategies can move out the risk spectrum, it is useful to think of three broad categories of aggressiveness, which I call "cautious all around", "modest reaching", and "pushing the limit". I think of these investor types being similar to Type I, Type II, and Type III skiers[16].

Cautious all around, Type I investors don't need to worry about fat tails with two exceptions:

- "End-of-the-world-as-we know it" scenarios which affect all investors regardless of how aggressive or conservative the investor is.
- Inadvertent exposure to fat tails, such as being invested in a "cash-plus" fund thought to be safe, but which turns out to have significant sub-prime exposure.

The primary risk faced by modestly-reaching Type II investors is that the risk premia they hope to harvest from their beta (market) exposures fail to materialize, and may even turn out negative ex-post. In addition, there is the question of what constitutes "modest reaching", and the possibility of reaching a little too far. This possibility can be avoided through broad diversification and scenario analysis to ensure adequate liquidity in a stressed environment.

Pushing-the-limit Type III investors need to have a risk-management function as well-developed as their investment function. The function may include non-normal modeling, hedging with non-linear derivatives, and an organizational structure designed to encourage creative thinking.

Please see Exhibit 7 (in appendix) for a comparison of these three investor types and the risk management issues they face.

V. Summary and Conclusions

1. Investing is concerned with the future, and the future is unknown. No model can change that.

2. Models are lenses through which we view aspects of the future. Models are by definition incomplete. Models are *tools* which have the potential to facilitate better investment decisions.

3. Portfolio theory is a useful way of thinking about future returns, diversification, and risk taking. As with any model, it has its limitations. It also has been subject to several misunderstandings which limit its acceptance, most notably the idea that standard deviation is synonymous with risk.

4. If one is over-reliant on portfolio theory – or on any model – one may be blind-sided by a fat tail or black swan. Put most simply, the risk management lessons of fat tails and black swans are:

[16] When you rent downhill skis in the United States, you are asked to classify yourself by how aggressively you ski. Cautious skiers and beginners are Type I, intermediate skiers are type II, and experts are type III.

- Fat tails: don't ignore history; the future may be like it.
- Black swans: don't ignore the possibility that the future may be very different than history

5. Risk is the possibility of something undesirable happening. Risk is not always measurable, but it can be a useful discipline to attempt to measure it, as long as one doesn't take one's measurements too seriously.
6. Risk management is the process of identifying what can go wrong with contemplated courses of action and taking steps to keep the likelihood and magnitude of bad outcomes within tolerable limits. Risk management is more challenging the more risk one takes, and may include liquidity and scenario analyses for moderately aggressive investors, and non-normal modeling and "tail insurance" strategies for more aggressive investors.
7. It can be very useful for investors to ask themselves, "What is the least risky thing I can do?" Risk-averse investors may decide to pursue the least risky option. More aggressive investors should also ask themselves this question, however, since it forces clarity in thinking and provides a useful benchmark against which additional risk and return can be measured.
8. If you want to see into the shadows, don't turn the light off. Shine another light from a different angle.

The events of 2008 have been painful for investors, and call for a fresh look at risk management and investment strategy. We at NEPC are learning along with our clients, and look forward to working with clients to put these learnings into practice.

References

Bhansali, Vineer (2008). "Tail Risk Management." *Journal of Portfolio Management*, Spring 2008.

Bookstaber, Richard (2007). A Demon of Our Own Design. Wiley.

Grantham, Jeremy (2009). Presentation to NEPC Offsite Meeting, January 2009.

Inker, Ben (2009). Presentation to CFA Asset Allocation Conference. March 2009.

Knight, Frank (1928). Risk, Uncertainty, and Profit.

Markowitz, Harry (1952). "Portfolio Selection." *Journal of Finance*.

Merton, Robert C. (2008). Address to Journal of Investment Management Conference, September 2008.

Rothschild, Michael, and Stiglitz, Joseph (1970.) "Increasing Risk: A Definition" *Journal of Economic Theory.*

Taleb, Nassim (2007). The Black Swan.

Exhibit 2

Illustration of Fat Tails and Black Swans with Balls and Urns

Fat tails: a simple example with known probabilities of extreme outcomes

Suppose you could choose one ball from one of two urns. Each urn contains a thousand balls. Depending on what color ball you draw from which urn, you get the following payoffs:

Table 1

Urn #1

1000 balls	payoff for drawing ball
100 are red	$40,000
800 are white	$50,000
100 are blue	$60,000

Urn #2

1000 balls	payoff for drawing ball
1 is red	*-$50,000*
998 are white	$50,000
1 is blue	$150,000

I have presented this hypothetical ball and urn problem to hundreds of students and colleagues over the years. The overwhelming majority choose a ball from urn #1, without hesitation. When asked why they chose urn #1 over urn #2, typically they say that the possibility of a significant loss if one draws from urn #2 makes that option less attractive, even though the two options have exactly the same expected payoff and standard deviation (see table 2). This illustrates that knowing the summary statistics of expected payoff and standard deviation isn't sufficient for knowing what really matters about the distributions[17].

Kurtosis is another summary statistic such as standard deviation or expected return[18]. It measures the extent to which a given standard deviation derives from:

- lots of small deviations (low kurtosis, no fat tails, urn #1), or
- a few large deviations (high kurtosis, fat tails, urn #2).

Like other summary statistics, kurtosis is useful to know if one does not know the actual distribution of possible outcomes and their relative likelihoods. Also, like other summary statistics, it is not of incre-

[17] In special cases, e.g. if one knows the distribution is normal, expected payoff and standard deviation are sufficient for knowing anything one wants to know about the distribution.

[18] Statisticians like to classify probability distributions by their "moments". The first moment is the mean – the average of all possibilities weighted by the likelihood of each possibility. Variance is the second moment, the expected squared deviation from the mean. Skewness is the third moment, the expected cubed deviation from the mean and also a measure of the asymmetry of a distribution. Kurtosis is the fourth moment, the expected fourth-powered deviation. The kurtosis of the normal distribution is 3. Excess kurtosis is the amount kurtosis exceeds three.

mental value to know a distribution's kurtosis if one already knows the distribution itself.

Let me illustrate. Which would you rather know about the urns?

- All of the information in table 1 and nothing else (that is, for each urn, the number of each color ball and the payoff associated with each color); or
- All of the information in table 2 and nothing else (that is, for each urn, the expected payoff, the standard deviation, and the kurtosis)?

Table 2

	Urn #1	Urn #2
Expected payoff	$50,000	$50,000
Standard deviation	$4,472	$4,472
Kurtosis	5	500
Excess kurtosis	2	497

Clearly, if you know everything about the urns, balls, and payoffs, there is no need to know the summary statistics. On the other hand, if you do not know the distributions and are given the summary statistics, the summary statistics are better than nothing. With means and standard deviations of the urn's payoffs' the same, we know to be wary of the distribution with high kurtosis, and would probably choose urn #1 on that basis. So knowledge of kurtosis helps us make the right choice in this scenario. However, this is a very circuitous way of making the decision, and completely unnecessary if one knows the underlying distributions.

To summarize thus far:

- Fat tails, also known as excess kurtosis, are present when the likelihood of extreme events is greater than is the case for a normal distribution with the same mean and standard deviation.
- When fat tails are present, knowing the mean and standard deviation of a distribution is not sufficient to evaluate the risk of the distribution.
- It is better to know the distribution than to know its summary statistics; if one knows the distribution, knowing the summary statistics does not provide additional information.
- If one does not know the distribution, knowing the summary statistics is better than nothing.

Black Swans: a simple example with unknown probabilities of extreme outcomes

I would like to change the facts of the ball and urn example slightly. Everything is the same in terms of the number and color of the balls in each urn, and with the payoffs. The only difference is that you do not know the number of each color in each urn. Instead, you learn this by sampling. At a cost of $10 per ball, you may draw as large a sample as you like from each urn prior to picking one ball for live payoffs. How many balls would you like to sample?

It is not unusual for subjects to want to sample 30 balls, thinking that that is the sample size that is sufficient for valid inference. Others will say, “hmm… 30 would makes sense if I knew the distribution was normal, but I don’t know this; why don’t I draw a hundred to be safe.”

So suppose you draw 100 balls from each urn. There is a good chance you will draw:

- Something close to 10 red, 80 white, and 10 blue balls from urn #1.
- 100 white balls from urn #2.

If all one knows about the proportions of colors in each urn is from these samples, most subjects now switch their live-payoff choice to urn #2, because it appears to be less risky. This illustrates the central idea of the black swan: the outcome you are trying to avoid – drawing a red ball from urn #2 – may not show itself in the data until it happens. This limits the usefulness of historical data in identifying and managing kurtosis, and puts a premium on:

- Understanding conceptually the sources of risk
- Forwarding-looking, imaginative analysis of what can go wrong

Exhibit 3

Deconstruction of Portfolio Theory

Characteristics of Portfolio Theory	Evaluation
Portfolio theory recognizes that: • The future is uncertain. • It is useful to distinguish between idiosyncratic risk and systematic risk. • Idiosyncratic risk is diversifiable. • Systematic risk is not diversifiable and therefore *warrants* a higher risk premium	All valid points.
Portfolio theory asserts that: • Diversification can reduce risk without reducing expected return • Systematic risk should be *expected to earn* a higher risk premium. • Leverage can expand the efficient frontier	Good rules of thumb but may not be universally true.
Portfolio theory models future returns with a multivariate normal distribution. This provides a system of equations which allows one to: • Calculate portfolio mean and variance given the means, variances, correlations and portfolio weights of individual assets. • Calculate the probability of any event, given the mean and variance of the distribution in question.	A useful approximation in many circumstances. However: • May underestimate the likelihood of extreme events (aka "fat tails"). • Does not distinguish between liquid and illiquid assets. If one believes fat tails and/or liquidity are important in a given context, it may be appropriate to complement portfolio theory with scenario analysis and liquidity analysis.
Portfolio theory uses the term "risk" to refer to the standard deviation of the probability distribution of future returns.	This causes a lot of confusion, but is just a semantic issue. Nothing about the multivariate normal model precludes using other risk measures. All of the various "downside" risk measures can be calculated within the context of the model.
Users of portfolio theory sometimes: • Confuse the statistical and English meanings of "expected return", resulting in target returns which have only a 50/50 chance of being realized • Over-rely on historical data in forming expectations for the future	These are potential implementation pitfalls; both are avoidable.

Exhibit 7

Fat Tail Management for Different Investor Types

	Type I	Type II	Type III
	"cautious all around"	"modest reaching"	"pushing the limit"
Description of investor type	Does something close to the lowest risk option; main bet is that the world as we know it will continue to exist	Reaches for higher return with mismatched beta exposures and modest alpha and illiquidity bets	Pushes the limits of financial engineering in the pursuit of superior risk-adjusted returns
What can go wrong?	1. inadvertently choose higher risk option than intended; 2. "end of the world as we know it" (which affects all investors, regardless of how conservative they are)	1. risk premia fail to materialize; 2. exposure to crisis event greater than planned, leading to liquidity problems. Also type I risks	1. Levering up on something that looks good in the model but in fact isn't. 2. Forced sales of illiquid assets. Also type I and II risks
Potential fat left tail	Small	Largely reflected in historical returns; primary additional risk is a potential liquidity squeeze	Risk of black swan is bigger when levering up on new ideas
Intensity of need for quantitative risk management	Small – traditional prudence sufficient	Medium – mean-variance analysis combined with liquidity analysis is appropriate	High – non-normal modeling and use of non-linear derivatives are appropriate
Need for qualitative risk management	Small/medium – traditional prudence sufficient; still important to ask, what can go wrong and how bad can it get?	Medium – important to think through the assumptions underlying expected risk premia	High – out of model thinking is an important check on models; organization needs to encourage imaginative thinking about what can go wrong