SHEDDING LIGHT ON THE FUTURE

ASSET ALLOCATION AND RISK MANAGEMENT IN A POST-CREDIT CRISIS WORLD*

Erik Knutzen, CFA, CAIA, Chief Investment Officer
John Minahan, CFA, Senior Lecturer in Finance, MIT

Introduction

The credit crisis of 2008 and subsequent episodes of market volatility remind us of our limited ability to peer into the future. When investors grapple with asset allocation they use various theories and models in an attempt to enhance their ability to forecast risks and returns to form views of the future. Yet many investors were caught off guard by the events of 2008, the "flash crash" of 2010, and the Euro zone debt crisis. Does this mean that our theories and models are flawed? Absolutely. Indeed, this is what it must be. No matter how good 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 sliver of the future, but the illumination also casts shadows. So, even though any one model, such as the greatly maligned but still useful mean-variance approach to portfolio construction—commonly known as modern portfolio theory—may have flaws, using multiple tools in a dynamic framework, as part of the critical asset allocation process, can increase luminosity, lighting up the future from multiple angles. In this paper we provide an overview of strategic asset allocation and risk management in the post-crisis world. We have three primary objectives:

– To review the challenges presented by non-normal distributions of outcomes, including tail events, to the traditional approach to asset allocation;

– To define risk management from first principles that is, principles general enough to encompass normal distributions, non-normal distributions, and no distribution at all; and,

– To describe a holistic approach to asset allocation and risk management in a dynamic framework using multiple models—we describe six, including mean variance—to dispel as many shadows as possible.

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

The most common approach when peering into the future is to assume that the aspects you care about—usually, risk and return for investors—can be represented by a probability distribution, that is, listing all possible values a variable can take on and assigning a 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 had repeated reminders that investment returns have neither of these properties. In fact, they display so-called fat tails and black swans.

To this end, a survey of the history of capital markets highlights that investment returns typically display extreme outcomes more often than a normal distribution with the same mean and variance would. This phenomenon can be quantified statistically with a measure called kurtosis; fat tails are referred to as excess kurtosis (Exhibit 1).[1]

Exhibit 1: Family of Distributions with Different Kurtosis

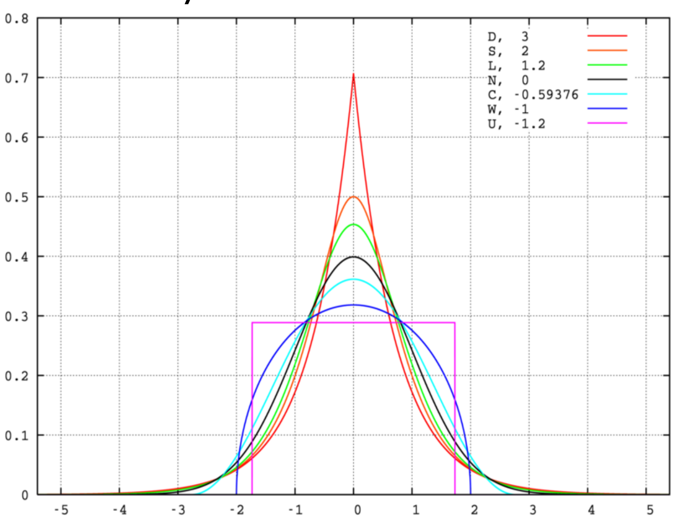

Source: Wikipedia

**This paper is an update and expansion of John Minahan's 2009 paper "Looking Into the Future Casts Shadows".*

[1] Both terms are derived 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 over this is called excess kurtosis.

Exhibit 2: An Assessment 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 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 it 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.

Source: NEPC

A much broader, related issue is the truism that *any* model of the future will emphasize some things and de-emphasize others and, will therefore, cast shadows that can create blind spots for users of the model. Sometimes problems can grow in these blind spots until they become big enough to draw attention to themselves, at which time they are experienced as so-called black swans, or unprecedented, hard-to-imagine, extreme events.[2, 3] 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 the Appendix for a detailed example illustrating the relationship between fat tails, black swans and kurtosis.

Greater awareness—intensified by the credit crisis—of fat tails and black swans has brought into question the relevance of normal distributions and the portfolio and capital market theories built upon them. In response, there are calls for new approaches to asset allocation and risk management that explicitly account for fat tails, black swans and downside risk.

That said, this doesn't necessarily indicate that mean variance portfolio theory is fatally flawed. Portfolio theory does indeed have shortcomings, but these are well known and can either be compensated for or avoided. Exhibit 2 deconstructs portfolio theory, laying out individual components to be assessed on their own merits.

To be sure, there are elements we need to unlearn about portfolio theory. We should stop thinking of standard deviation as a measure of risk. We should be aware that "expected" only

[2] 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; what had been an unappreciated difference between the model and reality now becomes a fully recognized black swan. This is similar to Bookstaber's (2007) idea that tight coupling combined with complexity will inevitably lead to an accident.

[3] The term "black swan" is derived from the fact that, for hundreds of years, all swans observed by Europeans were white. Philosophers used the idea of a black swan to illustrate the point that just because you haven't seen something, it doesn't mean it doesn't exist. Indeed, black swans did exist, and were eventually discovered by Europeans in Australia.

signifies a 50/50 chance of experiencing a target level. We need to be more explicit about portfolio theory being forward-looking and not dependent on historical data, though it is informed by it. And, in cases where fat tails or illiquidity may be present, we need ways of analyzing those risks. Most importantly, we have to recognize that portfolio theory is *just a model*. Like any model, it captures some elements of reality and misses others. Once we come to grips with these aspects of portfolio theory, we find it can serve as a very useful tool in a holistic approach to asset allocation and risk management.

What is Risk? And what it isn't.

Risk is the possibility of something undesirable happening.[4] It usually means the potential for poor returns in the context of an investment. Risk is not inherently measurable. It becomes measurable *to the specific investor in question* only by making assumptions about (a) an investor's ability to forecast probability distributions; and (b) how the investor weights the displeasure experienced from various bad outcomes. Unless one assumes that all investors are the same, there is no basis for defining a risk measure as characteristic of an investment as opposed to a characteristic of a specific investor's evaluation of that investment.

Risk is multi-faceted and, in part, specific to each investor. Important risks can include: risk of failing to meet an objective, risk of losing money (or drawdown risk), tail-risk or risk of extreme market outcomes, and liquidity risk. Also, risk can change through time.

But standard deviation is *not* risk. Unfortunately, the investment profession has fallen into some poor linguistic habits fueling enormous confusion around the mistaken idea that standard deviation and risk are synonymous. Some of this confusion may have contributed to what seems to be a popular revolt against portfolio theory. Portfolio theory is often criticized—mocked even—because it allegedly treats symmetrically upside and downside deviations from expectations. Yet, this is not a problem with portfolio theory per se but with the idea that standard deviation is risk, which is not central to portfolio theory.[5]

Denigrating the role of standard deviation in investment circles has become somewhat popular. Such a line of reasoning, however, ignores the fact that standard deviation can help quantify an element of risk that investors may find informative: the breadth of a distribution of potential outcomes. This, used in conjunction with additional models, can be a useful component of a comprehensive approach to risk management. Also, if you know the mean and assume the distribution is normal, you can calculate the probability of any possible future outcome. In fact, standard deviation can give a pretty good indication of other key risks such as the degree of a historical loss. Exhibit 3 compares historical standard deviation and maximum drawdown of major asset categories.

Exhibit 3: Volatility and Maximum Drawdown for Major Asset Classes

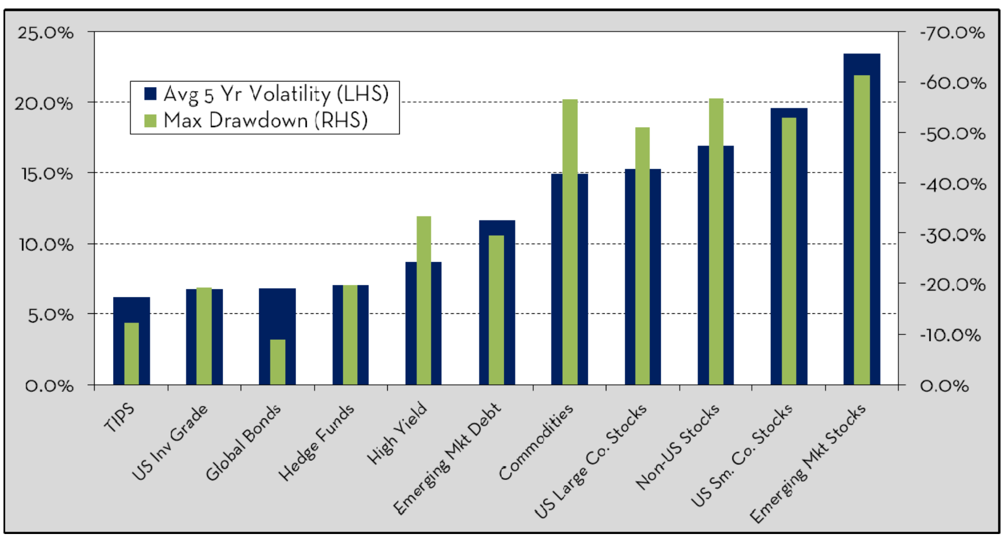

Source: Bloomberg, NEPC

Risk Management in the Presence of 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 magni-

[4] Knight (1921) 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.

[5] In the original paper on portfolio theory, Markowitz (1952) developed 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 (or standard deviation) with risk," came to be widely accepted as a cornerstone of portfolio theory, which, in turn, led to additional skepticism of the theory by practitioners who know the idea that risk is variance is silly. Yet nothing in Markowitz's theory requires that risk be equated with variance or standard deviation. His unfortunate comment to that effect, and the subsequent development of finance jargon conflating risk and standard deviation, seems to have led to the rejection of portfolio theory for the wrong reason.

tude of unacceptable outcomes within tolerable limits. Risk management is:

i) Awareness of the possibility of bad outcomes;

ii) Thinking through the potential consequences of actions;

iii) Avoiding courses of action that seem too risky; and

iv) Measurement.

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

1. What is the least risky course of action? 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. The least risky option could be cash, a liability-matching portfolio, a portfolio of Treasury Inflation-Protected Securities (TIPS), or anything that corresponds to an investor's way of thinking about risk minimization.

2. Does an investor choose to depart from the least risky course? Many investors choose to depart from the least risky course to seek a higher return. For instance, a higher level of risk may be chosen to pursue closing a funded gap in a pension program while minimizing contributions, or in efforts to grow spending rates in support of a charitable mission.

3. What risks are willing to be taken in an attempt to achieve performance above that of the least risky course? This acknowledges the investor's capacity for risk. The question also addresses how high the prospective returns need to be to justify the risk taken. It also seeks to gauge the willingness of the investor to withstand extreme outcomes and/or take on additional sources of risk such as illiquidity.

As investors move from lower risk to higher risk options, not only does risk increase, but also the ability to describe risk with a probability distribution decreases. This underscores the importance of using multiple models and a dynamic approach to asset allocation and risk management.[6, 7]

Exhibit 4: The Relationship between Risk and Standard Deviation

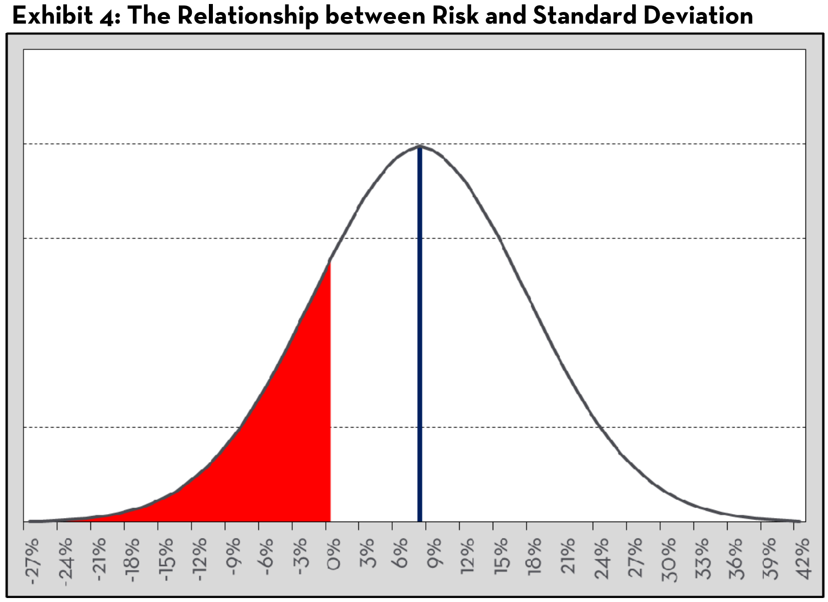

Source: NEPC

A Multi-Faceted Approach to Risk Management

The application of multiple models in the portfolio risk management process can help shine a spotlight on the future, while reducing blind spots. Below is a description of six models that can be used in a complementary fashion to assist with this important process. Looking into the future is challenging and no single, or multiple, model approach can provide all the answers. This observation also highlights the importance of applying qualitative and dynamic assessments to the process through time.

Mean Variance Analysis

In the classic approach to asset allocation, expected returns, standard deviations and correlations are developed for asset categories. Using mean-variance analysis, an expected return and standard deviation can be generated for any given asset mix, and with an optimizer, an "efficient frontier" can be created, describing an array of portfolios that represent the highest expected return for a given standard deviation (or vice versa). Despite the well-articulated short-comings of traditional portfolio theory, this approach can serve as a useful starting point in describing risk-

[6] These questions also highlight the importance of hedging known—but uncompensated—sources of risk. Examples of these may include developed market currency movements and, for defined benefit retirement plans, interest rate risk. These topics are addressed in greater length in NEPC white papers, including "Managing Developed Country Currency Risk - A Proactive Approach," and "Understanding Duration Risk in Pension Plans: The Case for LDI," available at www.nepc.com.

[7] This paper is focused primarily on addressing market-related risks, including liquidity risk and active management risk. Investors also face additional non-market risks such as those associated with operational issues and counterparties. A discussion of such operational risks is outside the scope of this paper.

Exhibit 5: Higher and Lower Standard Deviation Distributions

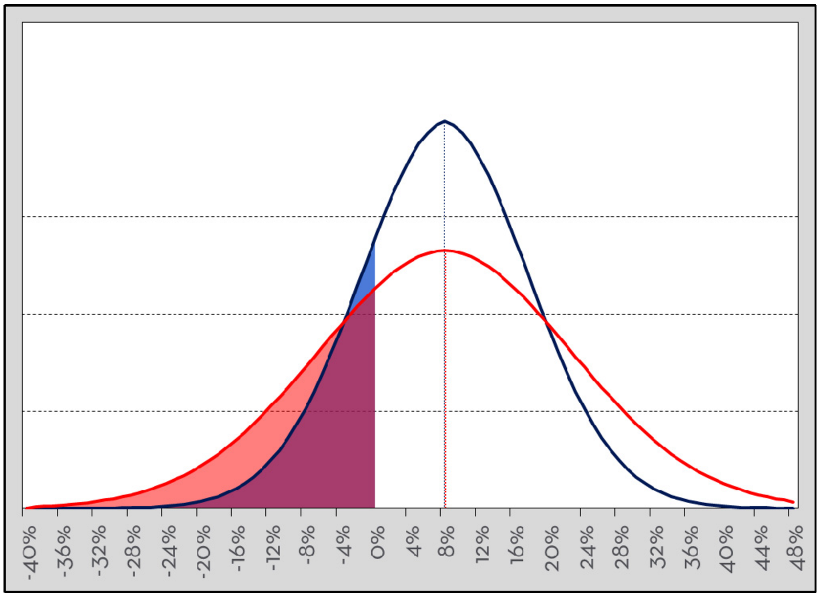

Source: NEPC

reward tradeoffs and diversification benefits. Indeed, standard deviation and risk are closely related (Exhibits 4 and 5). If one defines risk as the probability of a loss, then risk is, in part, determined by standard deviation. If one holds the mean constant, then an increase in standard deviation always implies an increase in the probability of loss.

Given the potential shortcomings of mean variance analysis, it is critical to use a forward-looking process using market pricing of key building blocks of expected return and risk to develop model inputs. It behooves investors to remember that "expected" only means that there is a 50/50 chance of experiencing a target level[8] because of the likelihood of divergence between modeled and actual outcomes. It is also important to view the analysis with statistical skepticism. The output of a mean-variance model is not prescriptive to the basis point, that is, there is little statistical difference between two portfolios with expected standard deviations of, say, 12%, but with expected returns that differ by, say, 0.20%.

Finally, the market environment of the last 30 years has created another challenge for users of mean variance analysis. Declining interest rates since the early 1980s have led to depressed forward-looking return expectations across markets, particularly for fixed income categories. In an effort to stretch for previously achieved-returns, many investors have chosen to move up the efficient frontier and take greater exposures to higher volatility asset classes such as equities. In order to get a handle on this approach to portfolio construction, it is important to incorporate additional tools.

Risk Budgeting

The traditional method of portfolio analysis involves allocating capital. In a risk budgeting analysis, risk is allocated to asset classes. This approach often uses the same asset class categories and inputs as traditional mean variance analysis, while examining contribution to risk by asset category. The most commonly cited example of an insight generated by risk budgeting analysis is regarding a 60% stock/40% bond portfolio. In the capital allocation framework, this common asset allocation appears diversified, yet a risk budgeting analysis reveals that as much as 90% of the expected volatility of the 60/40 portfolio will come from equities. In fact, this is not a diversified asset allocation, but a "one-bet" portfolio on equity risk as shown in Exhibit 6. With risk budgeting, an investor compares capital allocations—that is, how much of the portfolio's assets are allotted to each

Exhibit 6: Risk Budgeting Analysis of a 60% Stock/40% Bond Portfolio

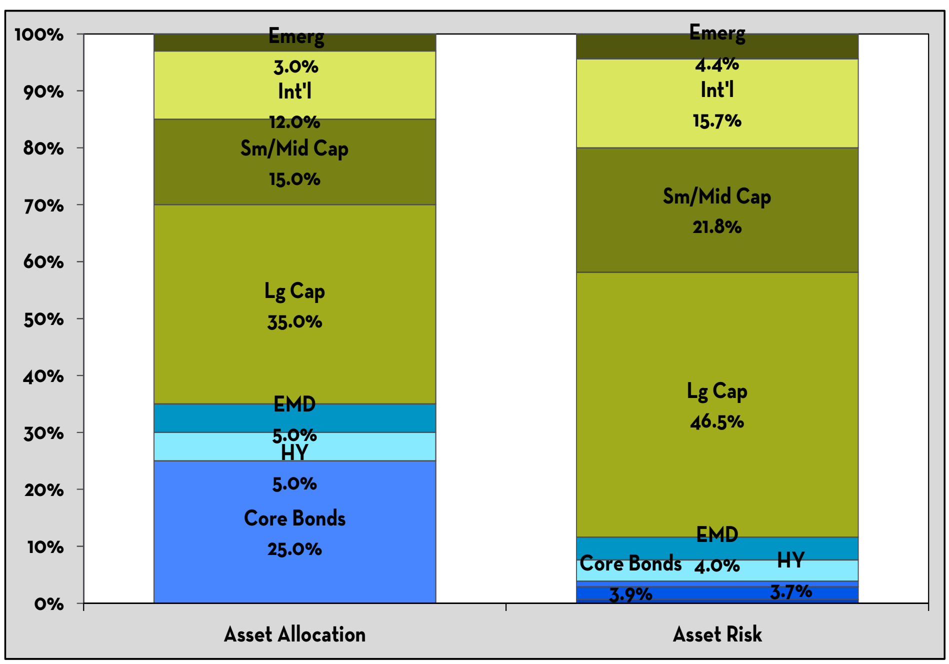

Source: NEPC (2013 five-to-seven years forecast for risk, return and correlations)

[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 return target, and ignores the extent to which the left tail expands as one reaches.*

category—to the risk from those investment categories.

For most investment programs, a risk budgeting analysis reveals that allocations to equity or equity-related investments dominate the program's risk budget even if they only represent 50%-to-60% of assets in a policy. Risk budgeting allows a better understanding of the sources of a portfolio's risks, and the relative benefits of diversification of other asset categories. For example, an allocation to commodities—a relatively volatile asset category in its own right—may reduce overall risk and consume very little of the program's risk budget due to low expected correlations to equities and other components of an asset policy.[9]

Risk budgeting, however, does rely on the same inputs as mean variance analysis, including assuming normal distributions of returns for all asset classes and stable correlations among those asset classes. Therefore, it is subject to some of the same shortcomings, including an inability to incorporate fat-tail events or black swans. As a result, more needs to be done.

Scenario Analysis

This approach entails modeling the effects of different economic environments on a portfolio. These environments can include recession, high inflation, stagflation, or more extreme outcomes, such as fat-tail events or black swans. Scenario analysis can help in asset allocation by allowing investors to examine how different asset mixes will perform in extreme environments and understand outcomes in the tails of return distributions. To this end, scenario analysis is a deterministic approach in contrast to the stochastic nature of mean variance and risk budgeting analyses which seek to describe a normally distributed and randomly generated set of outcomes.

If one is dealing with a known tail event, such as a particularly negative scenario, then risk measurement is straightforward. An investor can simply quantify the impact of the event and decide whether to accept an extreme negative outcome. In real life one almost never deals with a known tail event. At best we have historical experience and some conceptual models for affirming or modifying expectations about future outcomes. Nonetheless, it can be useful in scenario analysis to treat historical experience as if it is possibly representative of future outcomes.

Exhibit 7 shows an example of a scenario analysis using a base case outcome and a set of standard scenarios that include assumptions for how markets would behave in difficult environments. Through scenario analysis, an investor can ask whether they can withstand a certain extreme outcome by quantifying key variables such as asset value, funded status, or spending amount at the end of a period. Scenario analysis can also assist with assessing tradeoffs associated with efforts to avoid negative outcomes versus failing to capture upside. Investors can also test alternative asset allocation targets to check for a more desirable trade-off between negative and positive results, or, perhaps, consider applying specific hedges to limit extreme outcomes.

Exhibit 7: Sample Scenario Analysis

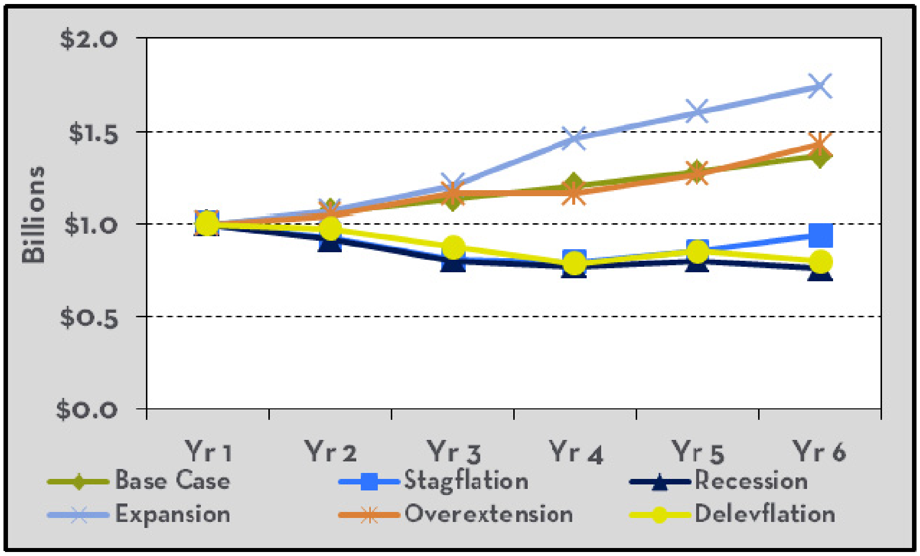

Source: NEPC

A Brief Diversion – historical return distributions, time diversification and tail-risk hedging

As an example of how one can use fat-tailed historical samples to inform future expectations, we examine the historical performance of US stocks. In the history of annual returns of the Standard & Poor's 500 Index (Exhibit 8), there was a fall of nearly 90% once in the last 85 years and declines of close to 40% occurred four times in the last 40 years. On the basis of this historical data, it would certainly be reasonable to assume that there is a remote probability of a 90% fall in the S&P 500 going forward, and a likely possibility that we will experience another 40% decline sometime in the

[9] *The use of risk budgeting in the context of corporate defined benefit pension plan management is discussed in NEPC's white paper "Risk Budgeting: A Focus on a Pension Plan's Biggest Risks," available at www.nepc.com.*

Exhibit 8: Annual Stock Market Returns Have Fat Tails (US Stocks 1870-2012)

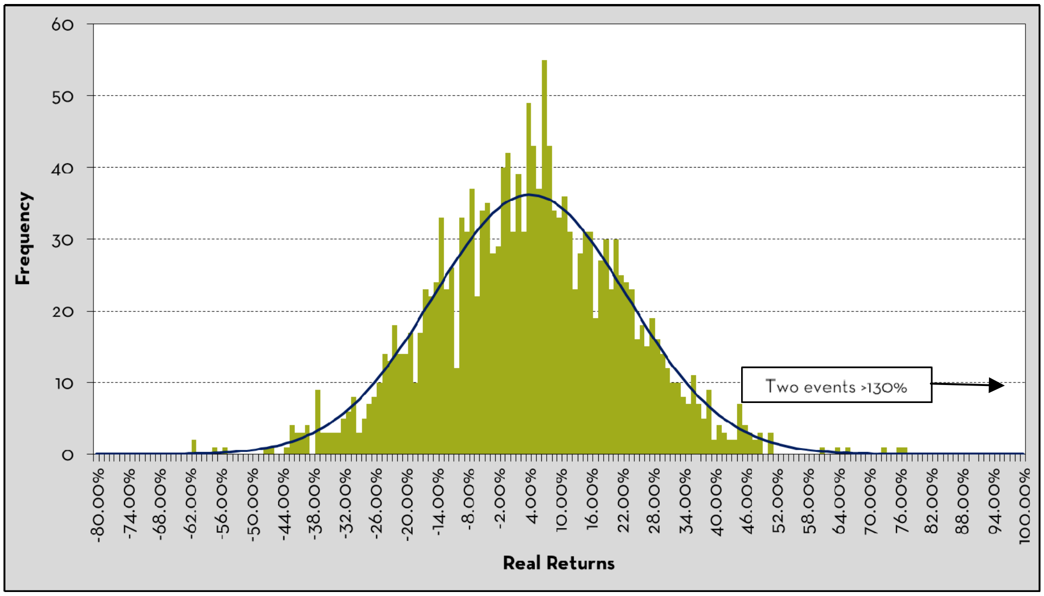

Source: Bloomberg

next decade or so. It is a matter of judgment which of these possible outcomes should be considered as a worst-case scenario. In this context, however, the stock market decline during the Credit Crisis does not qualify as a "black swan" as the result was contained in the historical return distribution; the results in the credit markets may have indeed been a black swan.

The idea, that return distributions become more normal-looking the longer the measurement horizon, is suggested by the central limit theorem and supported by a review of US stock market performance. Specifically, history shows 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 9).

If fat tails do decline with time horizon, then long-term investors can be less concerned with them as long as they do not become liquidity-constrained during the fat-tailed event. This has led to the observation that a long time horizon, combined with robust asset diversification, obviates the need to worry about tail events in the short-term. For long-term investors with ample liquidity, short-term market movements may create opportunities to buy assets at attractive prices. In fact, incurring the cost of hedging against tail outcomes for such an investor, by reducing the expected return of the program by the amount of the hedging cost, may amplify the risk of not achieving a target investment objective. This risk needs to be assessed carefully versus the benefits of the peace of mind from the short-term hedge. Instead, long-term investors with ample liquidity may be ideally positioned to be sellers of tail-risk insurance rather than buyers, an idea recently put forward by Robert Litterman.[10]

Liquidity Analysis

Balancing the requirements for and sources of liquidity is crucial to successful risk management. The first three models we have considered rely on changes in market prices—volatility or extreme outcomes—to define risk. Liquidity risk is a blind spot in these tools, and thus, separate analytics are required to inform investors how to size and manage this particular risk. In fact, an investor's time horizon effectively becomes very short-term if, during a fat-tailed event, he or she is forced to trade in an illiquid market. By selling illiquid as-

Exhibit 9: 30-Year Stock Market Returns Do Not Have Fat Tails (US Stocks 1870-2012)

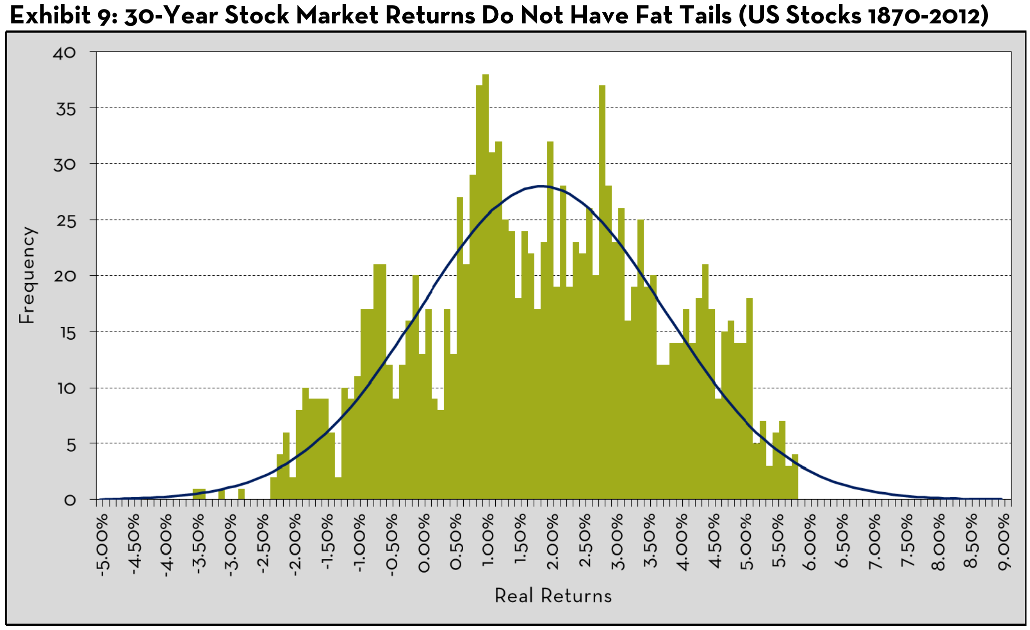

Source: Bloomberg

[10] *Litterman, Robert, "Who Should Hedge Tail Risk?" Editorial Financial Analysts Journal, May/June 2011*

sets at distressed prices, investors turn possibly temporary losses into permanent losses and give up the opportunity to recover from short-term dislocations.

Investors with a reasonably diversified portfolio will be primarily concerned with systematic fat tails, since idiosyncratic fat tails do not affect the total portfolio unless they are concentrated. One way of characterizing systematic fat tails is through the phrase "all correlations go to one in a crisis." Spiking correlations of risky assets is a short-term phenomenon; high correlations measured with monthly data do not imply that investments are tightly correlated over the long run. As a result, a dramatic rise in short-term correlations only creates a fat tail issue for investors forced to liquidate in the crisis.

The possibility of becoming a distressed seller during a fat tail event elevates forward-looking liquidity analysis in risk management and investment strategy selection, especially for investors with extensive private markets programs and/or who use leverage. The first step in this process is a detailed understanding of current portfolio liquidity, and known and potential calls on that liquidity. This information can then be combined with scenario analysis to "stress test" portfolios. Assumptions can be made on how the portfolio would behave in an extreme environment, including declining values of liquid asset categories, acceleration of capital calls, slowing of distributions, and/or increased demands for cash on the program.

Liquidity stress tests can identify the likelihood of becoming a distressed seller of illiquid assets to meet cash flow needs. Liquidity analysis can also offer insight into the "denominator effect" where, as public markets decline, the proportion of illiquid assets rises in the program's total asset mix. If this ratio gets too high, then a program may find itself "locked up" from an asset allocation standpoint without the latitude to correct imbalances in asset allocation or take advantage of new market opportunities with reallocated funds.

Economic Factor Analysis

The risk and return of different asset classes may be driven by common underlying factors such as economic growth, inflation and interest rates. Assessing the impact of such factors on a total investment program can provide deeper insight than more traditional categorizations of asset classes. A primary driver of equities is economic growth, which influences corporate earnings. Treasury bond prices move mainly due to changes in interest rates. Other asset categories may be influenced by more than one factor. For example, economic factor analysis indicates that high-yield bonds are heavily influenced by economic growth, more so than by changes in interest rates. According to this insight, high-yield bonds can be expected to experience price movements more in sync with equities than with Treasuries. Yet, in a traditional asset allocation modeling exercise, high-yield bonds would be grouped together with Treasuries into the fixed income bucket. Without considering the economic drivers of performance, a portfolio that appears diversified by asset class—using mean variance or risk budgeting analyses—may be dominated by common factor exposures.

Most long-term investment programs have very high exposures to economic growth, that is, the risk that economic growth disappoints, and correspondingly low exposures to less-correlated factors such as inflation and interest rates. Economic factor analysis can help investors assess the current biases of their program, whether those biases are desirable or whether they should be offset with exposures to other factors. Factor analysis can also focus questions on the exposure of a program's liabilities to such factors and shed light on their relationship to the asset exposures (Exhibit 10). For instance, a company retirement plan's liabilities will be primarily sensitive to interest rates (absent a cost of living adjustment feature in the benefit structure). As a result, the asset pool may be grouped into two components:

Exhibit 10: Risk Allocations by Economic Factor

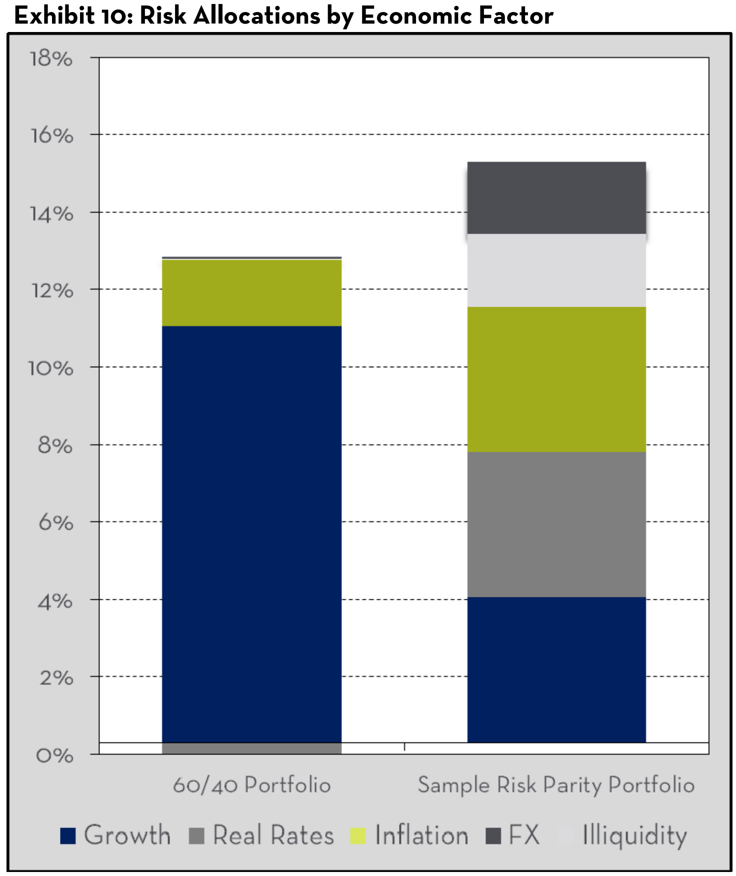

Source: NEPC

growth and an interest rate-hedge. In this case inflation may not represent a significant risk. Conversely, an endowment may be very sensitive to inflation, given the impact of rising prices on the ability of a charitable institution to fulfill its mission.

Academics and practitioners have identified a number of factors, including broad high-level drivers of security outcomes, such as economic growth, inflation and real interest rates, and more specific factors, for instance, country/region, sector/industry, momentum, value and quality. In applying economic factor analysis at a total investment program level—as opposed to within a single investment category such as equities—it is important to limit the number of factors used to those providing discrete information that is relatively uncorrelated with other factors. Our research indicates that five factors fit the criteria: economic growth, inflation, real interest rates, foreign exchange rates and liquidity.

USING MULTIPLE ASSET ALLOCATION MODELS IN A DYNAMIC FRAMEWORK CAN HELP REDUCE BLIND SPOTS

Even when focusing on a relatively limited number of factors, it is important to note that this type of analysis, while informative, is not definitive from a statistical standpoint. The correlations of economic growth, inflation and real interest rates with traditional asset classes are relatively low, although directionally informative. In fact, a comprehensive risk factor model, by definition, should encompass all the same information as a comprehensive asset class model.[11] Any difference between the outcome of a risk factor-based allocation approach and an asset class-based allocation approach represents a specific market view, or bet, relative to the global market portfolio. As a result, we caution against discarding traditional asset allocation tools and solely relying on economic factor analysis.

Active Risk Budgeting

Many investment program sponsors seek to increase returns by deploying active strategies, thus incurring an additional risk: the risk of failing to match a policy benchmark due to active management underperformance or negative alpha. One approach to mitigate this risk is applying an active risk budgeting model. This incorporates the volatility of active strategy returns relative to a benchmark or hurdle rate, and the correlations of these excess returns among the active strategies in the program. Using such a tool, an investor can understand the interactions of different active strategies and assess the incremental risk active managers contribute to a total investment program. This tool can also be used to evaluate whether a particular strategy's contribution to total active risk is consistent with the conviction level associated with the strategy.

A common insight from this analysis is that, due to the low correlations of many alpha strategies, the aggregate additional risk contributed to an investment program by active management is often quite modest relative to the risk stemming from exposures to asset classes or economic factors.[12]

A Qualitative and Dynamic Framework

One of the important lessons from the last five years is that key relationships among asset classes are constantly changing. The set-it-and-forget-it approach to asset allocation and risk management, where strategic policy is assessed every three-to-five years and inputs used are modeled over a 30-year horizon, is bound to develop significant blind spots. It is important to incorporate a dynamic approach to asset allocation combined with more qualitative assessments to portfolio positioning in light of market return and risk dynamics.[13]

Key elements of a dynamic asset allocation process include:

1. A more frequent review and adjustment of asset allocation using market-driven assumptions.

This allows investment program sponsors to adjust their strategic allocations as frequently as every year to take advantage of valuation opportunities, or to mitigate risks.

2. Incorporating an opportunistic component into asset allocation policy.

This enables investors to take advantage of severe and large market dislocations using an intermediate time horizon of two-to-three years. Such a dislocation occurred in the credit markets in

[11] Idzorek, Thomas M., Maciej Kowara, "Factor-Based Asset Allocation vs. Asset Class-Based Asset Allocation," Financial Analysts Journal, May/June 2013.

[12] For more on this topic please see the NEPC white paper "Applying a Risk Budgeting Approach to Active Portfolio Construction," available at *www.nepc.com*.

[13] For more on this topic please see the NEPC white paper "Investing in Volatile Times: A Dynamic Approach to Asset Allocation," available at *www.nepc.com*.

2008. Investors who made an opportunistic allocation to credit investment strategies at this time were able to capture significant additional return as those markets rebounded. More recent examples include distressed investing in Europe to take advantage of the Euro zone debt crisis, and pursuing direct lending strategies to replace traditional credit providers such as banks forced to curtail their activities due to regulation.

3. Delegating a portion of assets to flexible strategies such as global asset allocation and global macro.

RISK MANAGEMENT IS AN ONGOING PROCESS REQUIRING MULTIPLE TOOLS AND CONSTANT LEARNING

The last component of a dynamic approach recognizes that market prices fluctuate constantly. While most short-term changes represent noise, markets often move sufficiently away from fair-value for active managers to pursue profitable trades. A number of global flexible managers have shown that they have the ability to capture excess return or mitigate risk when included in long-term investment programs.

In addition to applying the multi-faceted risk tool framework described in the earlier section of this paper, investment programs may need to adjust their governance structures to implement a more dynamic approach to asset allocation. Changes may include expediting the investment decision-making process, delegating specific authorities to staff, or amending policy to include opportunistic and global flexible components of the strategic asset allocation. From a rebalancing standpoint, broadening policy bands can ensure that re-allocations are less mechanical and more flexible to reflect changing market relationships. Incorporating the output of the six models described above with an understanding of market risk and valuation can help investors pursue risk management within a dynamic framework.

Conclusion

The events of 2008 and the subsequent years have challenged investors. These experiences highlight the need for a fresh look at asset allocation and risk management. In fact, risk management is an ongoing process, requiring multiple tools and constant learning. New models need to be built and assessed. In addition, models need to be applied within a dynamic framework that can incorporate qualitative insights.

In summary:

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, by definition, are incomplete, but they 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 has been subject to several misunderstandings—most notably, the idea that standard deviation is synonymous with risk—that unduly limit its acceptance.

4. Over reliance on any one model can lead to being blind-sided by a fat tail or black swan. Don't ignore history; the future may be like it. Also, don't ignore the possibility that the future may be very different from history.

5. Risk is the possibility of something undesirable happening. Risk is not always measurable, but it is a useful discipline to attempt to measure it. Multiple models should be used for as much perspective as possible though one shouldn't take one's own 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. The more risk one takes, the more challenging risk management is. Therefore, additional models, such as scenario and liquidity analyses to stress test portfolios, may be beneficial.

7. It can be very useful for investors to ask themselves, "What is the least risky thing I can do?" This question can clarify thinking and provide a useful benchmark against which additional risk and return can be measured.

8. If you want to see into the shadows, don't turn off a light; shine another light from a different angle, and use as many sources of illumination as possible.

The authors thanks Mark Cintolo, Tim McCusker, Ian Spencer and Ben Swift for their assistance with this paper.

Bibliography:

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

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

Idzorek, Thomas M., Maciej Kowara. "Factor-Based Asset Allocation vs. Asset-Class-Based Asset Allocation." Financial Analysts Journal, May/June 2013

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

Knight, Frank (1921). Risk, Uncertainty, and Profit.

Litterman, Robert. "Who Should Hedge Tail Risk?" Editorial - Financial Analysts Journal, May/June 2011

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

NEPC white papers (available at *www.nepc.com*):

Dennen, Lynda, David W. Moore, "Risk Budgeting: A Focus on a Pension Plan's Biggest Risks." January 2010

Knutzen, Erik, "Investing in Volatile Times: A Dynamic Approach to Asset Allocation." December 2011.

McCusker, Tim, "Applying a Risk Budgeting Approach to Portfolio Construction." December 2010.

McCusker, Tim, "Managing Developed Country Currency Risk - A Proactive Approach." April 2011.

Moore, David W., "Understanding Duration Risk in Pension Plans: The Case for LDI." January 2010.

Appendix: 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

In this hypothetical ball and urn problem, most would choose a ball from urn #1, without hesitation. 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.

Kurtosis is another summary statistic such as standard deviation or expected return. 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 incremental value to know a distribution's kurtosis if one already knows the distribution itself.

To 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, you know to be wary of the distribution with high kurtosis, and would probably choose urn #1 on that basis. So knowledge of kurtosis helps 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

Now let's 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 to want to sample 30 balls, thinking that that is the sample size that is sufficient for valid inference. Another approach might be to 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

Disclaimers and Disclosures

255 STATE STREET, BOSTON, MA 02109 | TEL: 617.374.1300 | FAX: 617.374.1313 | www.nepc.com

BOSTON | ATLANTA | CHARLOTTE | CHICAGO | DETROIT | LAS VEGAS | SAN FRANCISCO