



OCTOBER 2014

BALLENTINE PARTNERS

Portfolio Optimization Theory versus Practice

“In theory there is no difference between theory and practice. In practice there is.”—Yogi Berra¹

This paper examines the gap between the theory of portfolio construction and its practice. In particular, it analyzes some of the problems in the application of portfolio optimization techniques to individual investors. We then discuss what can be done today to compensate for the problems with the theory, and what additional work needs to be accomplished.

Various books and articles have discussed some of the problems with Modern Portfolio Theory, which is the basis for most portfolio optimizers. This article compiles that information into a critique of the investment industry’s continued reliance on portfolio optimization tools.

Clients and investment professionals alike place far too much faith in the mathematics of portfolio optimization techniques and the recommendations that result from applying optimization techniques to clients’ portfolios. Clients will benefit from being more aware of the very substantial gap between asset allocation theory and the real world. Portfolio construction is mostly an art, not a science. Investment success depends upon experience, judgment, skill, and luck—the latter often showing up as fortuitous timing. Clients should avoid relying too much upon the seemingly precise construction of “optimum” portfolios and risk measurements produced by portfolio optimization tools.

What is portfolio optimization?

Nearly all investment professionals rely upon portfolio optimization techniques grounded in Modern Portfolio Theory (MPT) to structure investment portfolios for individual investors. Using statistical techniques and computer-assisted modeling, investment advisers are able to combine different types of assets such as stocks, bonds, cash, real estate, and hedge funds to create portfolios that claim to offer the best possible return for specified level of risk, or to minimize the amount of risk an investor must assume to achieve a specified amount of return.

The problem is that the real world does not always fit the theory perfectly. Sometimes, reality and theory are not even

close. Yogi Berra’s famous line about theory versus practice certainly applies to portfolio optimization techniques as practiced today. The remainder of this paper presents theoretical assumptions and generalities found in MPT with commentary from a “practice” perspective.

Problems with portfolio optimization models

We begin with an analysis of the assumptions underlying most portfolio optimization models. Following is a comparison of the gap between theory and practice.

Theory: The composition of each asset class is well defined and stable.

Practice: Asset classes are not well defined, and they are not stable.

There is no universally accepted definition of "asset class."² Private investors (and many investment professionals) who rely upon portfolio optimization models for assistance in portfolio construction typically ask very few questions about how asset classes are defined, how the composition of an asset class may have changed over time, and whether the historical data often used to estimate the future behavior of an asset class is appropriate for that purpose.

What actually distinguishes an asset class? There is no standard answer to the question. How many asset classes are required to describe the universe of equities—just one or several? Are hedge funds an asset class, or should they be considered a subset of global equities? There are no clear answers to these questions.

The composition of any asset class is not likely to remain stable for very long. For example, in the portfolio optimization process the equities of emerging markets are often represented by one of the emerging market indices. The composition of those indices has changed quite substantially over time. China and India are now among the world's largest economies, but they are still included in the emerging market indices. Is it reasonable to estimate the future behavior of China's economy based on the past results for an emerging markets index? Sometime soon China will no longer be indexed as an emerging market. The same sorts of questions could be asked about almost any asset class.

Theory: Each asset class can be represented by an index or some other metric that represents the performance of the asset class.

Practice: There is no universally accepted benchmark for any asset class. With some asset classes, investors face significant challenges to find data to support estimates of return, risk, and correlation. Unfortunately, some of the asset classes that seem to provide the most diversification benefits for investors are the asset classes for which investors have the least reliable performance data. For example, when private equity³ first captured the attention of individual investors there were only a handful of private equity funds. Investors were attracted to the high returns and diversification benefits that private

equity seemed to offer. Now there are many private equity firms and competition is fierce. Private equity benchmarks that rely upon historical returns for private equity may not be a fair representation of today's highly competitive private equity market. Venture capital⁴, hedge funds⁵, real estate⁶, and commodities⁷ suffer from similar performance measurement problems. These measurement problems present a very serious challenge in the portfolio construction process.

Theory: The only factors that matter in portfolio construction are the final portfolio's expected return and its risk parameters.

Practice: Investors certainly care about a portfolio's expected return and the relationship between risk and return. But, they also care about an array of other factors, none of which are taken into account by most portfolio optimization models, such as the following. This is where the financial planner can prove his or her value:

- How much annual cash flow does the investor require from the portfolio?
- How much income tax will the investor have to pay on investment income?
- How liquid is the portfolio, in case the investor needs access to immediate cash?

What is the unit purchase size for each asset class? The unit purchase size requirements for top notch hedge funds, venture capital, private equity, and private real estate funds may be beyond the reach of many private investors, including investors who are considered to be very wealthy.

- What are the diversification requirements for each asset class? For example, various studies have found that an investment in at least eight different funds with different strategies is required to achieve strategy diversification (Lhabitant and Learned, 2002). Diversification requirements for private equity, venture capital, and private real estate may require making new investments over a period of many years to achieve time diversification. In addition to strategy diversification, manager diversification must be considered.
- How much manager selection risk⁸ is associated with a particular asset class recommendation? Is the investor able to tolerate the consequences of a manager selection mistake?
- Over what time frame is the investor planning to

measure the portfolio results? Obviously, an investor who expects to live only a few years more and whose estate needs cash to pay estate taxes has a very different time frame than an investor who needs retirement income for 30 years or more.

- What are the consequences for the investor if the portfolio fails to perform as well as expected? An investor who needs funds on a specific date to pay a liability as it comes due faces different consequences than a wealthy grandparent who is investing to provide grandchildren with tuition assistance.

Theory: Investors are able to develop a sufficiently accurate forecast for each asset class that describes the mean return⁹, standard deviation of returns¹⁰, and the correlation of returns¹¹ between all relevant asset classes for any given time period.

Practice: It is highly doubtful that anyone can consistently and accurately forecast the above factors. Most forecasts rely heavily upon historical return data. If investment returns in each asset class have a tendency to revert to the mean, then a forecast based on historical data will tend to show a perverse bias favoring asset classes that may under-perform in the future.

The portfolio optimization process requires the following data elements as inputs for each asset class being considered for inclusion in the portfolio:

- Average expected return for each asset class
- Standard deviation of returns for each asset class
- Correlation of returns for each asset class with every other asset class being considered.

How many input variables is that? Assume that 10 asset classes are being considered. A portfolio optimization model requires:

- 10 estimates for average returns
- 10 estimates for standard deviations of returns
- 45 correlation factors

That's a total of 65 variables that must be correctly forecast.¹² What is the likelihood someone will be able to forecast returns, standard deviations, and correlations for multiple asset classes for a multi-year period with sufficient accuracy to produce a portfolio recommendation that is optimal, or nearly optimal? How much does it matter if the various elements of the forecast are inaccurate? It matters a lot,

particularly with respect to the expected return for an asset class. A relatively small change the expected return can cause the recommended investment portfolio to change quite substantially.

Theory: A rational investor would want to own a portfolio that lies on the (unconstrained) efficient frontier, because those portfolios will produce the highest return for any given level of risk. That is, the portfolios on the efficient frontier are the most desirable portfolios.

Practice: It is highly unlikely any rational individual investor would want to own a portfolio that lies on the efficient frontier.

Financial planning practitioners familiar with portfolio optimization techniques know portfolio recommendations produced by an unconstrained optimization are likely to be dominated by a small set of asset classes that happen to have favorable forecasts at that moment. If the return estimates are based on historical data, the recommended portfolio may be concentrated in asset classes that have performed well recently. It is likely at least some of the recently hot asset classes are now over-valued and about to perform poorly.

Would a rational investor prefer a highly concentrated portfolio thought to lie on the efficient frontier over a more broadly diversified portfolio that lies some distance from the efficient frontier? Perhaps, but only if the investor has a high level of confidence in the forecast for each asset class. In the author's experience, most investors have very little confidence in the forecast. Most advisers have so little confidence in their own forecasts they are unwilling to recommend the portfolios produced by the unconstrained optimization process.

One solution to this problem is to constrain the inputs to the portfolio optimization process so as to force the optimizer to produce a broadly diversified portfolio. The constraints usually include both a maximum or minimum percentage of the portfolio's capital allocated to each asset class. The constraints thus imposed change the location and the shape of the efficient frontier, so the recommended portfolio may then lie on the efficient frontier.

But, that is an admission the portfolio design process is mostly art wrapped in a veneer of scientific analysis. The adviser is forcing the optimization process to produce a preconceived

result. In the author's opinion, forcing the optimization process to produce a preconceived result is preferable to basing a recommendation on the unconstrained results of the optimization process. But, it is also an admission that, although the mathematics of portfolio optimization may be robust, the portfolio optimization process requires a liberal dose of experience and judgment.

Theory: The future pattern of returns for all asset classes can be accurately represented by a linear probability distribution, similar to the familiar bell curve¹³.

Practice: The distributions of returns for most asset classes do not follow the familiar pattern of the bell curve, or normal distribution. In a normal distribution, the probability an investment return will fall far above or below the mean—that is, more than three standard deviations from the mean, is very low.

Actual investment returns are far more likely to fall more than three standard deviations above or below the mean than would be predicted if actual returns were normally distributed (Mandelbrot and Hudson, 2004). The actual distribution of investment returns for most asset classes have “fat tails,” and are negatively skewed. The distribution of investment returns may follow some sort of exponential law, for which we presently do not have good modeling techniques. If the distribution of investment returns in the real world has “fat tails,” and is negatively skewed then the probability of extreme losses is much higher than is implied by standard risk models (see Figure 2).

Theory: Investment returns follow a random path. A market that is down in one period is equally likely to be either up or down in the next period, and vice versa.

Practice: The investment world is part of a dynamic economic system that has many very complicated feedback processes. Multiple extreme events occur within relatively short time intervals much more often than the theory suggests should be the case.

The pioneers who developed MPT understood that investment returns cannot be adequately described by a normal distribution. Later, in the 1960s, Benoit Mandelbrot (1963) suggested investment returns are best described by a Stable Pareto distribution (Mandelbrot, 1963). That type of

distribution has fatter tails than a normal distribution, meaning that the probability of extreme events is higher. The distribution can also be skewed. Working with a Stable Pareto distribution requires very highly specialized mathematical techniques. Interest in this problem died off until recently because the log-normal distributions commonly used for portfolio optimization seemed to do an adequate job. The stock market collapse of 2008 rattled the sense of complacency that built up around this problem.

Figure I: Distribution with fat tails

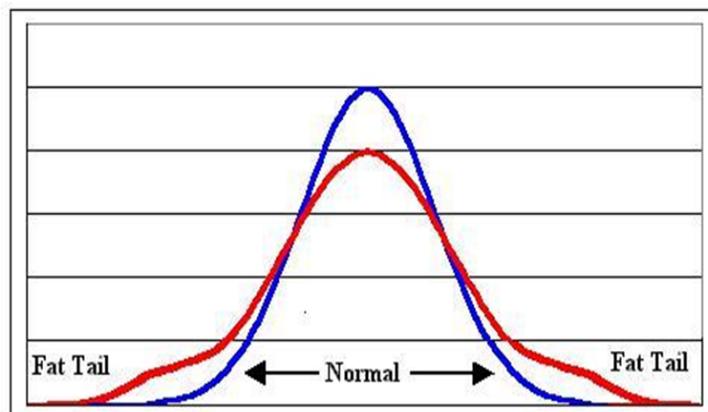


Figure I compares a normal linear distribution to a distribution with fat tails (theoryandreality.com/resources/rm_faq1.html). Although the differences between the two curves shown in Figure I may not look very significant, the amount of error introduced by assuming that investment returns can be described by a normal linear distribution is quite substantial. One consequence is the risk of loss in each asset class is significantly underestimated, and the risk of loss in the recommended portfolio may be considerably higher than is implied by the estimates produced by portfolio optimization models.

How far off are the risk estimates? According to risk estimates produced by portfolio optimization models that rely on normal distributions, market shocks and significant events like the one that began in mid-2007 and ended in March of 2009 are theoretically so improbable that they should never occur. A study by Cook Pine Capital in November 2008 summarized the problem as follows:

As an example, whereas the normal distribution of the daily return of the S&P would suggest a negative three-sigma event (-3.5 percent daily return) should have occurred 27 times over the last one hundred years, this has actually occurred 100 times in the 81 years since 1927. When one looks at even greater negative return days, the results become even

more pronounced. "...the "normal" likelihood of a negative four-sigma event (-4.7 percent daily return) is once in one hundred years; yet we have seen this take place 43 times since 1927. The same normal distribution suggests virtually no possibility (.00003 percent) of a day where negative returns are greater than 5.8 percent, but, once again, we have witnessed such days on 40 occasions in the last 81 years, and alarmingly, three times in 2008 alone. (pages 5 – 6)

In Figure 2, the red line shows the distribution of events predicted by a normal distribution. The blue bars show the actual frequency of occurrence of events. Clearly, the actual returns show many more severely negative events than are predicted by a normal distribution.

Theory: Portfolio optimization models provide accurate and concise estimates of the risk of loss in a portfolio. Most models estimate the risk of loss to two decimal places.

Practice: The models are incapable of providing this information with sufficient accuracy to be useful to investors who must make risk management decisions. Optimizations fail on two counts. First, investors have no way to estimate the probability of rare, but potentially catastrophic, events. Second, we have no way to estimate how catastrophic a rare event will be. This is a severe problem for those engaged in risk management (Taleb, Nassim. 2007). (For a thought-provoking discussion of the problem of estimation error and the difficulty of assigning probabilities to rare events, see Taleb, Nassim Taleb's article, "Black Swans and the Domains of Statistics.")¹⁴

Theory: Investment returns follow a random path. A market that is down in one period is equally likely to be either up or down in the next period, and vice versa.

Practice: The investment world is part of a dynamic economic system that has many very complicated feedback processes. Multiple extreme events occur within relatively short time intervals much more often than the theory suggests should be the case. However, it is very unlikely that an investor will be able to use this information to consistently make money by betting on extreme events. Between January 1, 1929 and September 13, 2012, an analysis of the daily returns of the S&P 500 Index shows there were 86 days when that daily return was 4 or more standard deviations below the mean, and 73 days when the daily return was 4 or more standard deviations above the mean. 80 days are about 0.4% of the total days during that period.¹⁵

For example, Figure 3 shows actual returns for the S&P 500 Price Index, and the years in which extreme market losses occurred. The Cook Pine Capital paper defined an extreme event as a single day market loss that was worse than -4.6 percent, which is equivalent to an occurrence that is four or more standard deviations below the mean, if investment returns are assumed to be normally distributed. The historical results for the S&P 500 Price Index contain many more of these extreme events than are predicted by theory (based on a normal distribution). Notice too that the extreme events were clustered together. In other words, it appears that extreme losses tend to occur in close temporal proximity to one another.

Figure 2: Actual S&P 500 Daily Returns for S&P Price Index vs. Normal Distribution¹⁶

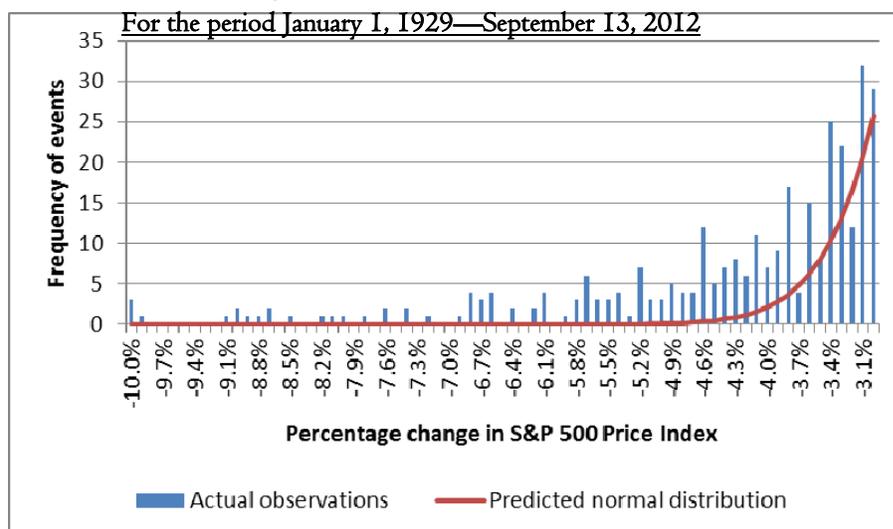
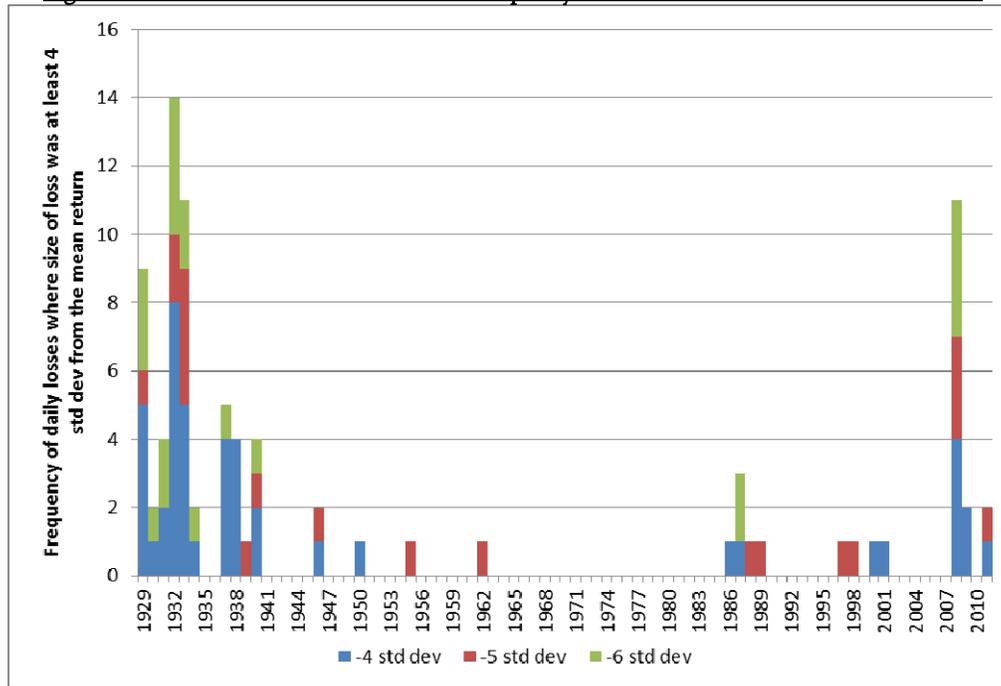


Figure 3: S&P 500 Index Price Returns—Frequency of Extreme Loss Events 1929 - 2011¹⁸



Theory: The standard deviation of returns for each asset class, and for the portfolio as a whole, is an appropriate way for investors to measure risk.

Practice: There are multiple problems with this part of the theory. For example:

- As explained above, any model that relies on standard deviation to describe the distribution of returns may be substantially underestimating the frequency of extreme events.
- Using standard deviation as a measure of risk implies that any deviation from the mean return is bad—including returns that are higher than the mean return. In reality, investors are worried about losing money; they have no problem if they make more money than they were expecting.¹⁷
- Few private investors have a great appreciation for standard deviation as a measure of risk.
- Most portfolio construction techniques use the annual standard deviation of returns for each asset class, but it is more probable that investors tend to react to short-term volatility. As explained in the following paragraph, for most investors the amount of short term volatility they experience is much worse than the annual standard deviation of returns seems to imply.

To understand just how much the annual standard deviation of returns masks the amount of volatility that investors actually experience, consider the following statistics using the

S&P 500 Total Return Index (SPX) as an example. In Table 2 and Figure 4, the calendar year return measures the change in the S&P 500 Total Return Index from January 1 to December 31. The intra-year maximum decline is the maximum percentage loss an investor experienced during the year, measured between whatever two points during the year produced the largest loss, while always moving forward in time. Thus, the time period associated with that loss varies from year to year.

The intra-year maximum increase is the maximum percentage gain during the year, measured using a methodology equivalent to the methodology used to determine the Maximum Decline. The average annual standard deviation of returns was 18.5 percent. It is easy to see the range covered by the Maximum Decline to Maximum Increase during most years far exceeded one average standard deviation. Small wonder that some individual investors have trouble sleeping at night.

Theory: The expected return, standard deviation, and correlations between asset classes remain stable over time.

Practice: None of these factors are stable. Nearly all portfolio optimizers are very sensitive to changes in expected returns, standard deviations, and correlations between asset class returns. Investment professionals know all of these factors tend to be unstable. For example, Figure 5 shows the five-year rolling average of the standard deviation of returns for U.S. large cap stocks for the period January 1876 to August 2012.

Correlations between asset class returns are also unstable. For example, Figure 6 shows the five-year rolling correlation between the returns for U.S. large company stocks and U.S. long term government bonds for the period from 1930 through August 2012. Investors need to be aware that the correlations between asset class returns is constantly changing, and may substantially deviate from the long term trend line for many decades.

One of the keys to managing investment risk is to own a portfolio of assets whose investment returns are independent of one another. In other words, if some asset classes are losing money investors hope to have other asset classes in the portfolio that are showing gains, or at least that all asset classes are not losing money at the same time. That is the goal of diversification. Unfortunately, under conditions of severe market stress, such as occurred in the second half of 2007 and continued through 2008, many investors who thought their portfolios were well-diversified were surprised that diversification failed to protect them from losses.

Table 2: S&P 500 Total Return Index 1990—2011¹⁹ - Calendar year returns, maximum percentage gain from low during year, and maximum percentage loss from high during year

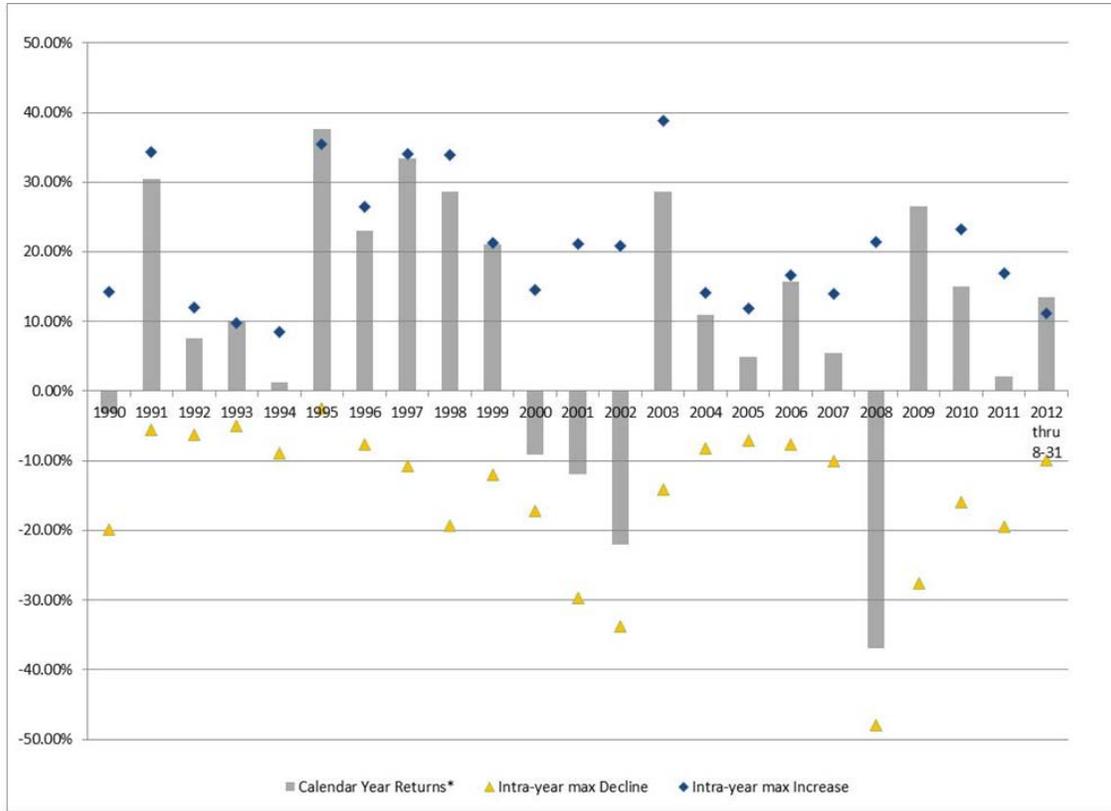
Year	Calendar Year Returns*	Intra-year max Decline	Intra-year max Increase
1990	-3.13%	-19.92%	14.23%
1991	30.00%	-5.60%	34.27%
1992	7.43%	-6.24%	11.93%
1993	9.92%	-4.99%	9.76%
1994	1.28%	-8.94%	8.46%
1995	37.12%	-2.53%	35.41%
1996	22.68%	-7.64%	26.49%
1997	33.10%	-10.80%	34.05%
1998	28.34%	-19.34%	33.86%
1999	20.89%	-12.08%	21.21%
2000	-9.03%	-17.20%	14.56%
2001	-11.85%	-29.70%	21.18%
2002	-21.97%	-33.75%	20.87%
2003	28.36%	-14.05%	38.86%
2004	10.74%	-8.16%	14.14%
2005	4.83%	-7.17%	11.89%
2006	15.61%	-7.70%	16.62%
2007	5.48%	-10.09%	13.90%
2008	-36.55%	-48.01%	21.36%
2009	25.93%	-27.62%	66.70%
2010	14.82%	-15.99%	23.20%
2011	1.57%	-19.39%	16.91%

Figure 7 shows price volatility and co-movement of prices for some commonly used asset classes. The asset classes included in the statistics for Figure 7 include: U.S. large capitalization equities, developed international equities, emerging market equities, real estate, commodities and U.S. government bonds. Figure 7 shows that periods of high market volatility have tended to coincide with periods when the prices of those assets classes have moved in the same direction at about the same time. In other words, during periods of high volatility diversification failed—at least with respect to those commonly used asset classes. Put another way, during periods of high volatility all these asset classes moved together up and down. But, it is only the coincidence of losses that poses a problem for investors.

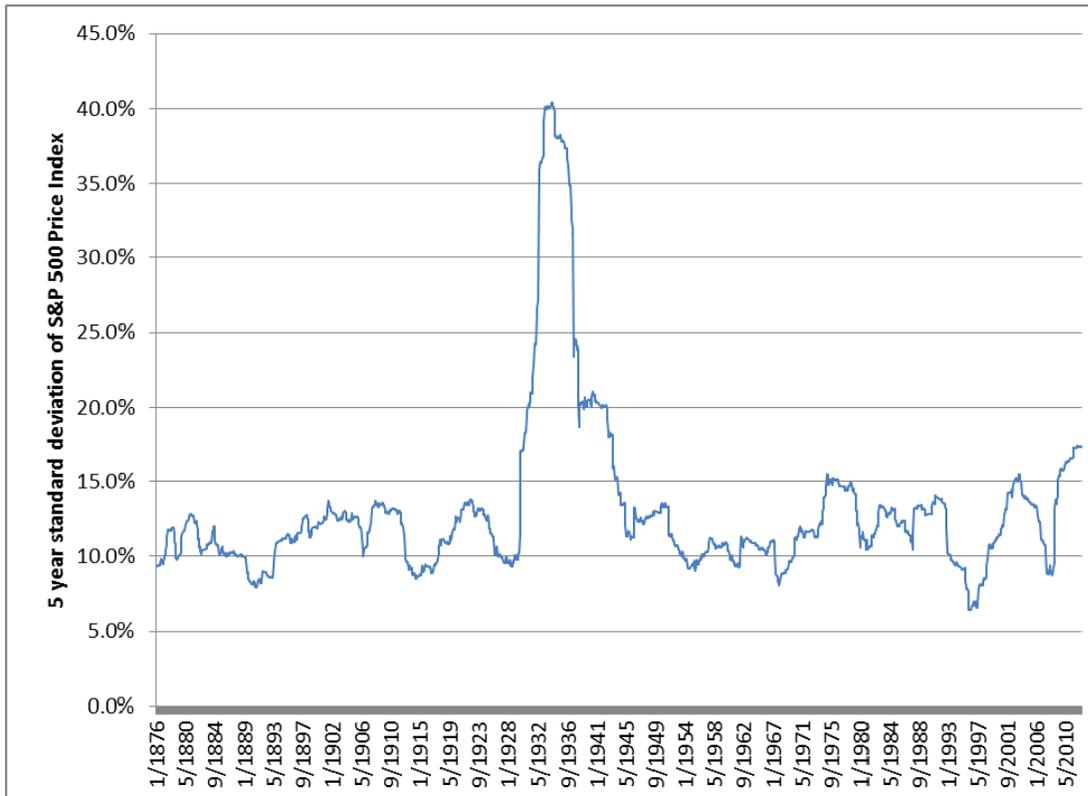
Theory: Investors are indifferent about the liquidity²⁰ of their investment portfolios.

Practice: Most private investors have a very strong preference for liquid investments, and they demand compensation in the form of substantially higher expected returns when making an illiquid investment (McCulloch, J. 1975).

Figure 4: S&P 500 Total Return Index 1990—2011 21— Calendar year returns, maximum percentage gain from low during year, and maximum percentage loss from high during year



**Figure 5: 5-Year Rolling Standard Deviation of Returns for Large U.S. Stocks²²
January 1876 through August 2012**



In practice, the liquidity problem is addressed in some models by applying an adjustment to the costs or returns of illiquid assets to make them appear to be less attractive. More often, investors or their adviser will simply place constraints on the amount the optimization tool is allowed to allocate to illiquid investments. The constraints are necessary because illiquid asset classes, such as private equity, venture capital, hedge funds, and private real estate funds, are often depicted as having relatively high expected returns and low correlations with other asset classes. Those features appear to make the illiquid asset classes very attractive as portfolio diversifiers. If illiquid assets are not constrained, the mathematics of portfolio optimization may result in a portfolio dominated by illiquid investments. Ironically, individual investors have often been disappointed in the performance of their illiquid assets because those investments have failed to perform as projected. The problems can be traced back to the issues discussed earlier; namely, inability to gain access to top-performing managers, reliance on historical data to project future returns, inadequate diversification, the vintage year problem, etc.

Theory: Any asset class represented in the optimization process can be purchased in any size, without loss of efficiency.

Practice: Private investors face substantial obstacles to obtain efficient access to high quality investments in many asset classes. This is particularly true of hedge funds, private equity, venture capital, commodities, and private real estate—the asset classes which portfolio optimization models depict as being very important to constructing a diversified portfolio. The problems include, for example, inability to gain access to top-performing managers, high minimum investment amounts, and lack of availability of suitable and attractive investment opportunities when the investor is ready to deploy capital.

Figure 6: Correlation of Returns for Large U.S. Stocks and Long-Term U.S. Government Bonds²³

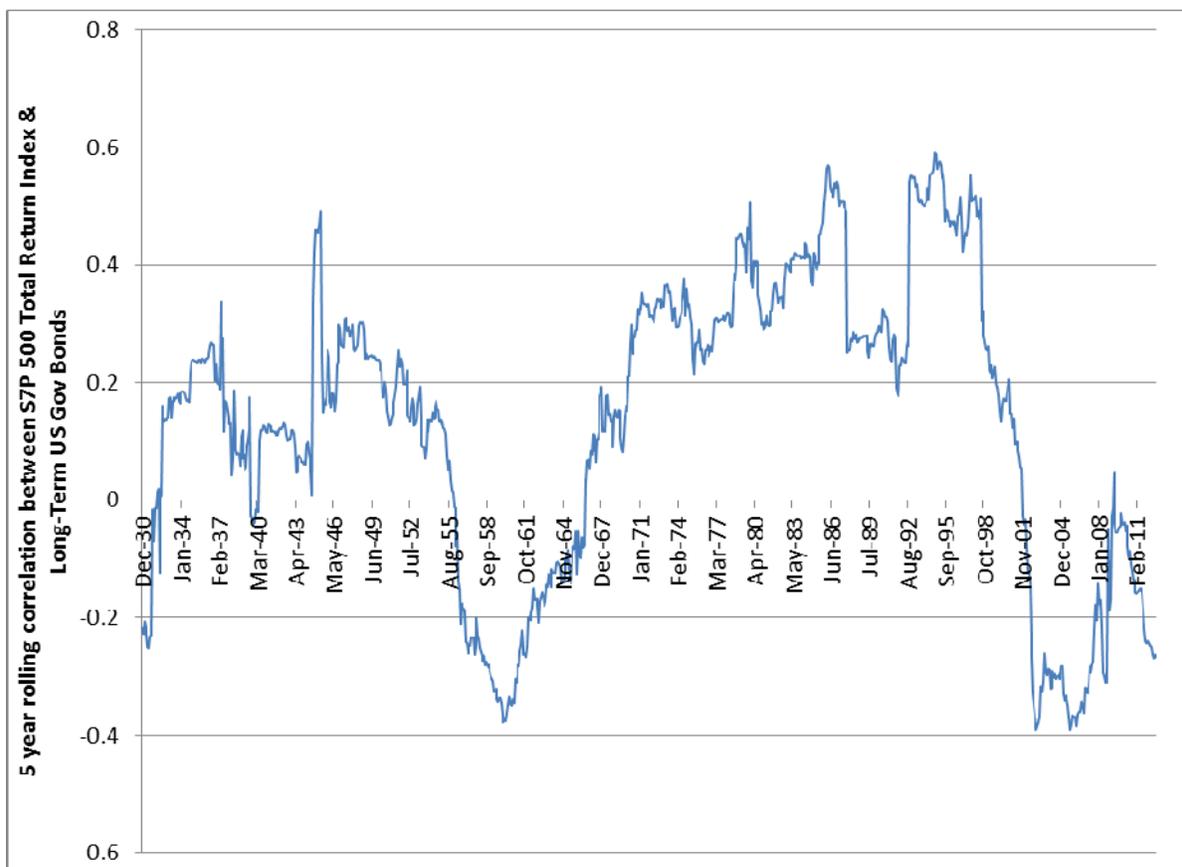
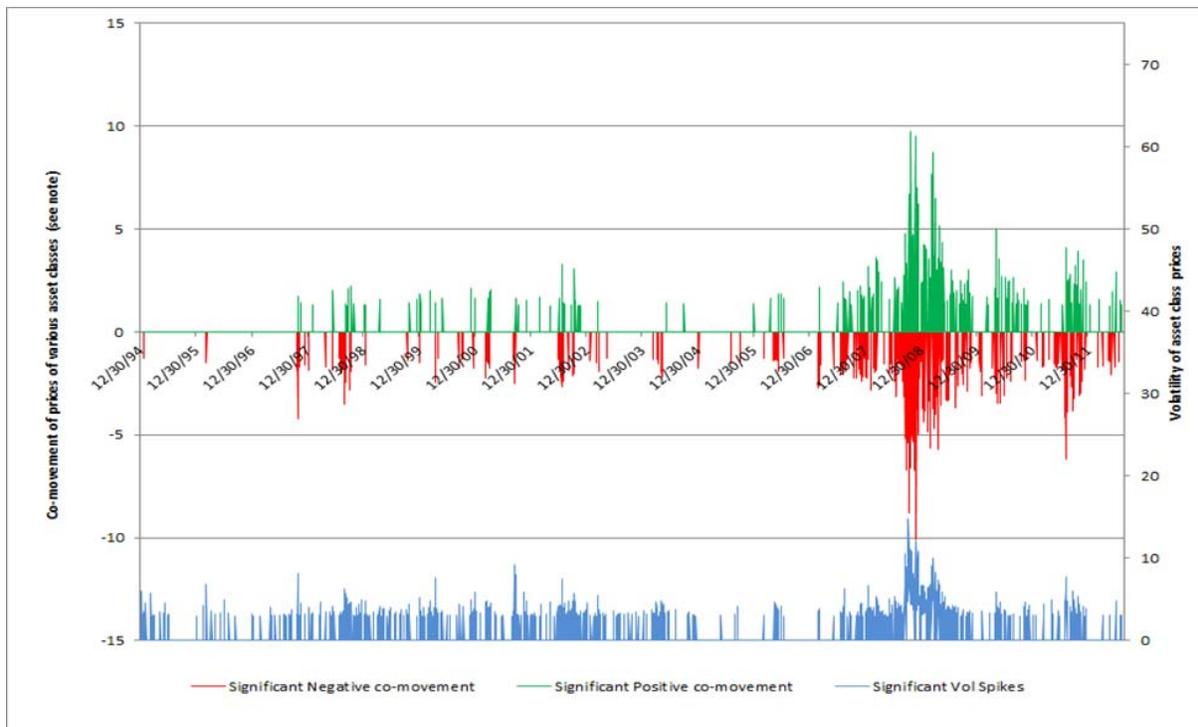


Figure 7: Market Volatility and Coincident Price Movements of Commonly Used Asset Classes²⁴
January 1, 1995—June 30, 2012



Theory: Investors are able to easily rebalance their portfolios at least once a year by moving capital between asset classes to maintain the desired asset allocation.

Practice: Private investors face significant costs, timing problems, and cash flow problems when they seek to rebalance in non-tax-deferred accounts. Investors incur taxes and transaction costs that may cause the investor to indefinitely postpone rebalancing. Annual rebalancing is effectively impossible for many taxable investors within some illiquid asset classes. That portion of the portfolio can be rebalanced only when a liquidity event occurs that frees up capital. Investors typically have no control over the timing of liquidity events in their illiquid investments.

Theory: The performance of the specific investments selected to implement the asset allocation strategy will be closely aligned with the assumptions used to develop the strategy.

Practice: This gap between theory and practice is due to the way advisors decide to implement the portfolio recommendations developed with the optimization model. The implementation method selected by many investment firms is not well aligned with the assumptions underlying the strategic asset allocation plan. Many investment firms fill their clients’

portfolios with products whose historical and projected performance bears little resemblance to the performance of the asset class index used during the portfolio optimization process. Many firms recommend active managers whose activities trigger substantial fees, trading costs, and taxes that are not reflected in the index that was used during the portfolio optimization process. The risk that a manager may under-perform relative to the index is also not accounted for. Investment advisers can mitigate this risk by using an appropriate mix of index investment vehicles and active managers.

Where does the future hold?

It is obvious the portfolio optimization tools widely used in the investment world rely upon assumptions that do not fit particularly well in the real world. Yet, these tools are widely used and relied upon by individual investors and advisors as if the underlying assumptions are valid.

The first step is to acknowledge that asset allocation and portfolio construction is mostly an art, not a science. Advisers have an obligation to inform clients of the limitations of the theory. Clients will benefit from being more aware of the very substantial gap between asset allocation theory and performance in the real world. Investment success depends upon

judgment, skill, and luck—the latter often showing up as fortuitous timing.

Financial planners, investment advisors, and clients should avoid relying upon the seemingly precise construction of “optimum” portfolios and the risk measurements produced by the application of portfolio optimization tools. For example, some optimization programs produce an estimate of the “worst case” portfolio loss based upon a negative annual return three standard deviations from the mean expected return. Such risk measurements are likely to substantially underestimate both the frequency and the severity of losses. Actually, history has shown the “worst case” that can and does occur in the real world is probably far worse than the imagination of even a very capable adviser is able to conjure up.

Portfolio optimization tools are useful to assist planners, advisers, and clients in ways to conceptualize the asset allocation problem. But, they should not be relied upon as the primary tool for structuring an investment portfolio. A more reasoned approach is to construct a portfolio around the client’s objectives, informed by the use of statistical tools such as those included in portfolio optimization models. This approach requires a disciplined process that is able to incorporate subjective factors into the analysis. The subjective factors may include a wide range of issues that are important to the client, including cash flow requirements, risk tolerance, and need for liquidity. This is a more fitting approach for investors than the artificial endeavor of striving to place a portfolio on a theoretical “efficient frontier.”

Clients and portfolio managers following such an approach should ask themselves the following questions:

- What measures of risk make sense given the objectives?

It is reasonable to assume that the majority of individual investors are primarily concerned with achieving at least a minimum target return that is necessary to meet their financial goals, while minimizing the risk of significant losses—that is, investors tend to be risk averse. Very few investors understand standard deviation as a measure of risk. Rather than looking at the projected standard deviation of their portfolio, investors and advisers should seek to understand how much the portfolio loss might be under different kinds of extreme conditions. This can be assessed by running simulations, such as Monte Carlo simulations or scenario simulations using historical data from various time periods.

- How much does the client have at risk, realistically, and can they afford those risks?

This type of information can be obtained through scenario analyses using historical data. However, clients and advisers should be mindful that the future may reveal even worse scenarios than the typical American has ever experienced.

- What is the investment outlook for each major asset class that will be included in the portfolio?

Answering this question requires that the advisor and client have an opinion about the relative pricing of asset classes at the time the portfolio strategy is being formed. Just because a trend is expected (higher inflation) doesn’t mean an asset class, such as commodities, will necessarily perform well if it is already overpriced in the market.

- What asset classes make sense in the portfolio?

The assets selected should make sense based on the client’s objectives, tax situation, liquidity requirements, risk policy, and the investment environment that is expected. Investments should be grouped into an asset class based upon an expectation that they will perform in a similar manner, but differently from other asset classes, when it matters most—that is, under conditions of extreme stress in the financial markets. This analysis requires careful thought. Trying to create asset classes based upon statistical analysis is not likely to lead to valid conclusions because there is too little data available about extreme events.

For example, a client may believe that she has achieved diversification by putting both emerging market and domestic equities into her portfolio, because on most days these asset classes perform differently. However, in more extreme environments, when diversification is most needed to meet the client’s risk management objectives, these classes of assets might actually behave quite similarly. Common sense would predict this. A common sense approach will lead to fewer asset classes being included in the portfolio because there are only a few categories of assets that behave differently under extreme conditions.

- What alternative investment scenarios might occur, and how will the asset allocation fare in these scenarios?

Investors should “stress test” their strategies against various environments, asking how each asset class might perform under those circumstances, and how the portfolio as a whole will fare. The client and the adviser should probe for

underlying connections between seemingly diverse asset classes to see if there is some common factor that could cause diversification to fail. To do this, the advisor and client may need to collaborate to explore risk factors lurking in the client's situation. This is especially important for clients who are business owners or investment professionals.

- What will drive investment selection?

Once a portfolio strategy is established, a key goal is to ensure that each asset class in the portfolio adheres to its assigned mission. Finding managers who may outperform their benchmarks should be secondary. Often, this step gets muddled by investors and advisers, who may be lured by performance. A classic example of this is a fixed income manager taking credit risks to beat the benchmark, when the manager's role in the portfolio is to provide downside market protection and to assure access to liquidity, if needed.

Advisers and clients also need to incorporate the financial planning aspects of portfolio construction when designing an investment portfolio. This is hard work, which many investment advisers are not well-equipped to do because they lack the time, resources, and training necessary to do it well. However this work is essential to help the adviser develop a deep understanding of the client's risk tolerance, cash flow requirements, return objectives, liquidity requirements, and tax situation. Specific tasks include:

- Carefully analyzing the client's balance sheet to identify concentrations of risk, determine available liquidity, and to examine the relationships between assets and liabilities;
- Assessing the income tax and transaction costs associated with changing the client's current portfolio holdings, including:
 - How much will it cost to make changes in the portfolio? (It is important to take into account that the

client must pay these costs up front before any of the benefits of portfolio restructuring are realized.) How long will it be before the client recovers these costs?

- What is the risk that the costs may not be recovered within an acceptable timeframe?
- Analyzing the client's cash expenditures from today through the end of life expectancy and through estate settlement²⁵;
- Analyzing the client's reliable sources of cash flow from earnings, dividends, interest, real estate income, etc.;
- Netting reliable cash income against cash requirements to determine the expected amount and timing of cash flow deficits over the entire forecast period;
- Developing a plan to fund the cash deficits from the investment portfolio by matching assets with liabilities (the liabilities being the cash flow deficits);
- Measuring the portfolio withdrawal percentage in each year that has a cash flow deficit. Withdrawal percentages that consistently exceed 4 percent are probably not sustainable for the long term (Cooley, Philip)²⁶ ; and
- Using scenario analysis to determine how large a cash reserve the clients need to have to allow them to feel comfortable sticking with their strategy during challenging investment periods. For example, what if the client suddenly lost his or her job?

The search is already on for more useful portfolio optimization tools, both within the field of financial planning and in colleges and universities. Until a model is developed that can incorporate the shortcomings of today's MPT models, it is important for financial planners to remember that no statistically based tool can adequately substitute for a disciplined and thoughtful approach that considers the practical constraints of the real world and the real needs of investors.

References

¹ Some sources attribute this to Jan van de Snepscheut, a Dutch Mathematician.

² The CFA Institute has published a recommendation about how asset classes should be defined. However, the CFA definition is not universally accepted.

³ Private equity funds invest in the equity and convertible debt of private companies, and they sometimes conduct buyouts of public companies that result in a delisting of public equity.

⁵ As used in this paper, “hedge fund” refers to a pooled investment vehicle that uses investment strategies that are typically not permitted in registered mutual funds, such as borrowing, shorting, and extensive use of derivative contracts.

⁶ As used in this paper, “real estate” includes both publicly listed Real Estate Investment Trusts (REITs) and private funds that invest in various kinds of real estate.

⁷ As used in this paper, “commodities” refers to publicly listed funds and private funds that invest primarily in energy, agricultural products, industrial metals, and other raw materials.

⁸ Manager selection risk is the risk associated with selecting the wrong manager for an asset that requires active management. If an asset class requires active management, then the investment results may depend heavily upon the skill of the manager. Manager selection risk is very high for hedge funds, private equity, private real estate, and venture capital funds. In other asset classes, public equities for example, investors have a choice between passive investments with little or no manager selection risk versus active management with higher manager selection risk.

⁹ The mean return is the arithmetical average of the distribution of returns.

¹⁰ The standard deviation of returns is a statistical factor that measures the variability of returns for a given time period. For example, the historical returns for stocks have a much higher standard deviation than the returns for bonds for most time periods. This reflects the fact that the prices for stocks tend to fluctuate much more than the prices for bonds. MPT uses standard deviation as a measure of risk. According to Modern Portfolio Theory, asset classes with high standard deviations are riskier than asset classes with lower standard deviations.

¹¹ A correlation is a statistical factor that describes how one variable moves in relationship to another. Correlations can be any value between -1 and +1. If two variables tend to move in the same direction at the same time—both going up or down, then they are positively correlated. If they move in opposite directions—one going up while the other is going down, then they are negatively correlated. Correlations describe only the relationship between two variables. If the relationship is not linear, the correlation calculation may produce an erroneous result. This is a problem for MPT optimization programs because the relationships between investment returns for different asset classes may not be linear.

¹² Some portfolio optimization models use factor analysis to reduce the number of variables that must be forecast. Factor analysis is a statistical method used to describe variability among observed, correlated variables in terms of a potentially lower number of unobserved variables called factors. It is possible, for example, that variations in three or four observed variables mainly reflect the variations in fewer unobserved variables. Factor analysis searches for such joint variations in response to unobserved latent variables. The information gained about the interdependencies between observed variables can be used later to reduce the set of variables in a dataset (see http://en.wikipedia.org/wiki/Factor_analysis).

¹³ Most portfolio optimization models use a log-normal distribution, which is a modified version of the bell curve that is associated with the normal distribution. A linear probability distribution can be described by a linear function (i.e. a function with no exponential factors). Recent research suggests that investment returns may be distributed according to a fractal or power function, that is, the distribution of returns may be non-linear.

¹⁴ Taleb, Nassim. 2007. “Black Swans and the Domains of Statistics.” *The American Statistician* 61, 3 (August)

¹⁵ A version of this chart appeared in Study of Fat Tail Risk, by Cook Pine Capital, Nov 26, 2008, available on the Internet at: <http://www.cookpinecapital.com/pdf/Study%20of%20Fat-tail%20Risk.pdf>. The chart was updated for this paper. S&P 500 Price Index data were used to compute the daily return for that index for the period January 1, 1929 to September 13, 2012. We then computed the average daily return and measured the deviation of returns for each day from the average return. We grouped the daily returns according to where they fell relative to the mean. For example, if the return for a particular day fell two standard deviations from the mean return, that day was counted as a two standard deviation “event.”

¹⁶ *Ibid.*

¹⁷ Practitioners have tried to address some of the shortcomings of MPT by using measures such as semi-variance to focus on the risk of loss. That is an improvement, but it still leaves other serious problems with the theory.

¹⁸ Please refer to note 14. Here is an example of how to read Figure 3. In 1932, the S&P 500 Price Index experienced 8 days with losses that were -4 standard deviations below the mean return, 2 days with losses that were -5 standard deviations below the mean, and 4 days with losses that were at least -6 standard deviations below the mean.

¹⁹ Ballentine Partners, LLC chart prepared by Katerina Barka using data from Bloomberg and Yahoo Finance for the S&P 500 Total Return Index (SPX). The intra-year maximum percentage increase was measured by selecting the highest value that produced the largest percentage increase in value, relative to a low point that occurred earlier in the year. The maximum intra-year decline was computed in a similar manner.

²⁰ “Liquid” in this context means readily convertible into cash under a broad range of market conditions and at a price that can be easily determined in advance of the transaction.

²¹ Ibid.

²² Ballentine Partners, LLC analysis prepared by Gregory Peterson, Ph.D., CFA based on Robert Shiller’s data. Five-year rolling standard deviation of S&P Price Index. We used the Price Index because most investors experience market volatility primarily through price changes, rather than changes in total return.

²³ Ballentine Partners, LLC analysis prepared by Gregory Peterson, Ph.D., CFA based on Dimensional Fund Advisors’ data. Five-year rolling serial correlation between S&P Total Return Index and Long Term Government Bond Index. The chart and an explanation is available at www.ballentinepartners.com under “Press & updates, White papers.”

²⁴ Ballentine Partners, LLC analysis prepared by Gregory Peterson, Ph.D., CFA based on Bloomberg data. The data are the daily returns of the MSCI US, MSCI EAFE, MSCI EM, Bloomberg North American REIT, Thompson Reuters Commodity and Barclays Intermediate Treasury Index. The chart shows principal component analysis. (For an explanation of principal component analysis, see: http://en.wikipedia.org/wiki/Principal_components_analysis and see Johnson, Richard A. and Wichern, Dean W., “Applied Multivariate Statistical Analysis,” third edition, Prentice Hall, 1992.) The positive and negative co-movements of prices are measured on the left vertical axis. The co-movement of prices is analogous to the standard deviation generalized to co-movements of prices across multiple asset classes. The high volatility periods are measured on the right vertical axis and are analogous to standard deviation in the context of multiple random variables. The sensitivity threshold for co-movement of prices to count as an “event” that is shown on Figure 7 was $\geq +1.28$ for positive movements, and ≤ -1.28 for negative movements. The sensitivity threshold for volatility of prices was 2.92. Those thresholds screen out all but the most extreme 20% of observations when a spike in volatility was accompanied by the co-movement of asset prices either up or down. The chart and an explanation is available at www.ballentinepartners.com under “Press & updates, White papers.”

²⁵ This analysis calls for the creation of various scenarios, including both spouses living well beyond actuarial life expectancy (50% of the population is still alive at life expectancy), the primary income-earning spouse dying early, disability of either spouse, loss of income by the primary income-earning spouse, etc.

²⁶ Cooley, Philip, et al. “Sustainable Withdrawal Rates From Your Retirement Portfolio”: www.afcpe.org/assets/pdf/vol1014.pdf.

Lhabitant, F., and M. Learned. 2002. “Hedge Fund Diversification: How Much is Enough?” FAME Research Working Paper 52 (July): <http://ssrn.com/abstract=322400> or <http://dx.doi.org/10.2139/ssrn.322400>.

Mandelbrot, B. and R. Hudson. 2004. *The (Mis)behavior of Markets*. New York, NY: Basic Books.

Mandelbrot, Benoit. 1963. “The Variation of Certain Speculative Prices.” *Journal of Business* 36 (October): 394–419.

Rivera, Peter. “Theory & Reality Risk Management.” www.theoryandreality.com/resources/rm_faqI.html

Taleb, Nassim. 2007. “Black Swans and the Domains of Statistics.” *The American Statistician* 61, 3 (August)

McCulloch, J. 1975. “An Estimate of the Liquidity Premium.” *Journal of Political Economy*, 83: 1.

Cooley, Philip, et al. “Sustainable Withdrawal Rates From Your Retirement Portfolio”: www.afcpe.org/assets/pdf/vol1014.pdf.