



ASSET ALLOCATION

Windham Investment Review

Mark Kritzman

Windham Capital Management, LLC
5 Revere Street
Cambridge, MA 02138
www.windhamcapital.com
617.576.7360

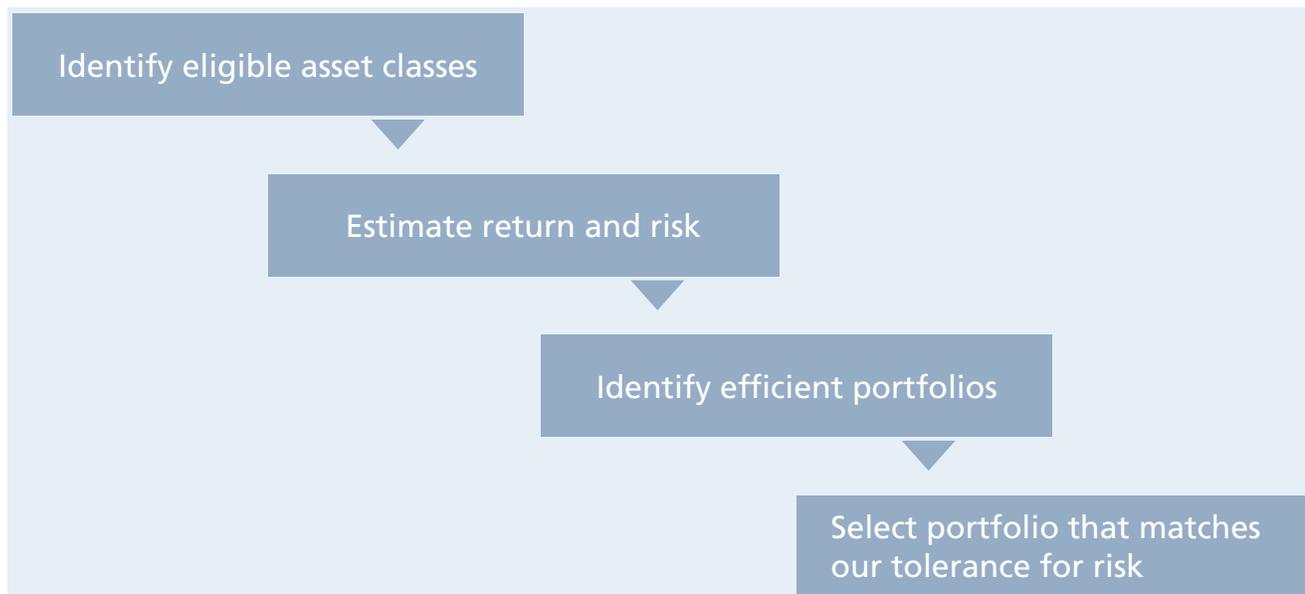
November, 2007

ASSET ALLOCATION

Asset allocation is one of the most important and difficult challenges we face as investors. Thanks to Harry Markowitz, we have an elegant and widely accepted theory to guide us, though implementation in the face of real world complexities is less straightforward than theory might suggest. In this inaugural issue of the *Windham Investment Review* we describe how to determine allocation to broad asset classes given the complexities of the real world.

There are four steps to asset allocation (Figure 1). We must first identify eligible asset classes. Then we need to estimate their expected returns, volatilities, and correlations. Next we must isolate the subset of efficient portfolios that offer the highest expected returns for different levels of risk. Finally we need to select the specific portfolio that matches our tolerance for risk.

Figure 1: Four steps to asset allocation



01 Eligible Asset Classes

What constitutes an asset class? First, we should expect an asset class to improve our portfolio's efficiency either by raising its expected return or by lowering its risk. Consider commodities, for example. We might believe that their expected return is insufficient to raise our portfolio's expected return because advances in technology tend to outpace depletion of resources, thereby lowering commodity prices. However, because commodities offer diversification against financial assets, especially during periods of high unanticipated inflation, their inclusion in our portfolio might lower risk thus offsetting their expected reduction of return.

We should also require homogeneity among the components of an asset class so that we do not forego opportunities for diversification. If an asset class comprises dissimilar components, then by investing in it we implicitly impose the unnecessary and potentially harmful constraint that the components must be held in the same relative proportions as their weights in the asset class. We should be able to achieve a more efficiently diversified portfolio if we partition the dissimilar components into multiple asset classes.

Finally, an asset class should be sufficiently large to absorb a meaningful fraction of our portfolio. If we were to invest in an asset class with inadequate capacity, we would likely drive up the cost of investment and reduce our portfolio's liquidity. The consequence might be to lower our portfolio's expected return and increase its risk to the point at which the proposed asset class would no longer improve our portfolio's efficiency. For illustrative purposes we will consider the following asset classes in our asset allocation analysis. Keep in mind, though, that these are but a few of the many other asset classes available to us.

Asset Classes:

U.S. stocks

Foreign stocks

U.S. bonds

Real estate

Commodities

Cash equivalents

02

Estimating Expected Returns, Standard Deviations and Correlations

The standard approach to asset allocation is based on portfolio theory, which requires us to estimate expected returns, standard deviations, and correlations.¹ To estimate expected returns, we start by assuming markets are fairly priced; therefore, expected returns represent fair compensation for the degree of risk each asset class contributes to a broadly diversified market portfolio. These returns are called equilibrium returns, and we estimate them by first calculating the beta of each asset class with respect to a broad market portfolio based on historical standard deviations and correlations. Then we estimate the expected return for the market portfolio and the risk-free return. We calculate the equilibrium return of each asset class as the risk-free return plus its beta times the excess return of the market portfolio. Admittedly, the markets are seldom, if ever, in equilibrium, but the pull in this direction is powerful and persistent. Moreover, we can easily adjust the expected return of each asset class to accord with our views about departures from fair value.

Suppose we estimate the market's expected return to equal 7.5% and the risk-free return to equal 4.0%.² Given these estimates, together with estimates of beta based on monthly returns from January 1997 through December 2006, we derive the equilibrium returns shown in Table 1.

Table 1: Expected Return

Asset Classes	β	Equilibrium	Views	Confidence	Blend
U.S. stocks	1.65	9.85%	9.85%	NA	9.85%
Foreign stocks	1.70	10.02%	10.02%	NA	10.02%
U.S. bonds	0.30	5.04%	6.00%	100%	6.00%
Real estate	0.86	7.01%	7.50%	50%	7.25%
Commodities	0.30	5.05%	6.00%	50%	5.52%
Cash equivalents	0.00	4.00%	4.00%	NA	4.00%

¹See, for example, H. Markowitz, "Portfolio Selection," *Journal of Finance*, March 1952. In 1990 Markowitz was awarded the Nobel Prize in Economics for his development of portfolio theory.

²The current risk-free return is readily observable. There are a variety of methods for estimating the expected return of a diversified market portfolio. For example, we might adjust the historical risk premium to accord with current risk levels and add this adjusted risk premium to the current risk-free return

We may expect some asset classes to produce returns that differ from those that would occur if markets were in equilibrium and perfectly integrated, especially if they are not typically arbitrated against other asset classes. Suppose we expect real estate to return 7.5% and commodities to return 6.0% and that we assign as much confidence to these views as we do their equilibrium returns. We can blend these views to derive our expected returns. The final column of Table 1 shows the expected returns for each of the asset classes in our analysis.

We also need to estimate the standard deviations of the asset classes as well as the correlations between each pair of asset classes. These values, shown in Tables 2 and 3, are estimated from monthly returns for the period beginning in January 1977 and ending in December 2006.

It is important to note that standard deviations and correlations are not always stable through time. It is therefore useful to separate historical returns into those returns associated with normal times and those associated with periods of market turbulence.³ This separation allows us to estimate these values for each regime and to stress test portfolios by measuring exposure to loss based on risk characteristics that prevail during turbulent periods.

Table 2 shows estimates for standard deviations for both normal and turbulent regimes, while Table 3 shows the correlations for both regimes. As you might suspect, volatility rises during times of turbulence.

Table 2: Investment Options

Asset Classes	Normal	Turbulent
U.S. stocks	16.69%	22.99%
Foreign stocks	18.75%	24.46%
U.S. bonds	5.83%	8.99%
Real estate	15.22%	22.26%
Commodities	17.40%	24.18%
Cash equivalents	0.95%	1.41%

³See, for example, Chow, G.,E. Jacquier, M. Kritzman, and K. Lowrey, "Optimal Portfolios in Good Times and Bad," *Financial Analyst Journal*, May/June 1999.

Table 3: Correlations

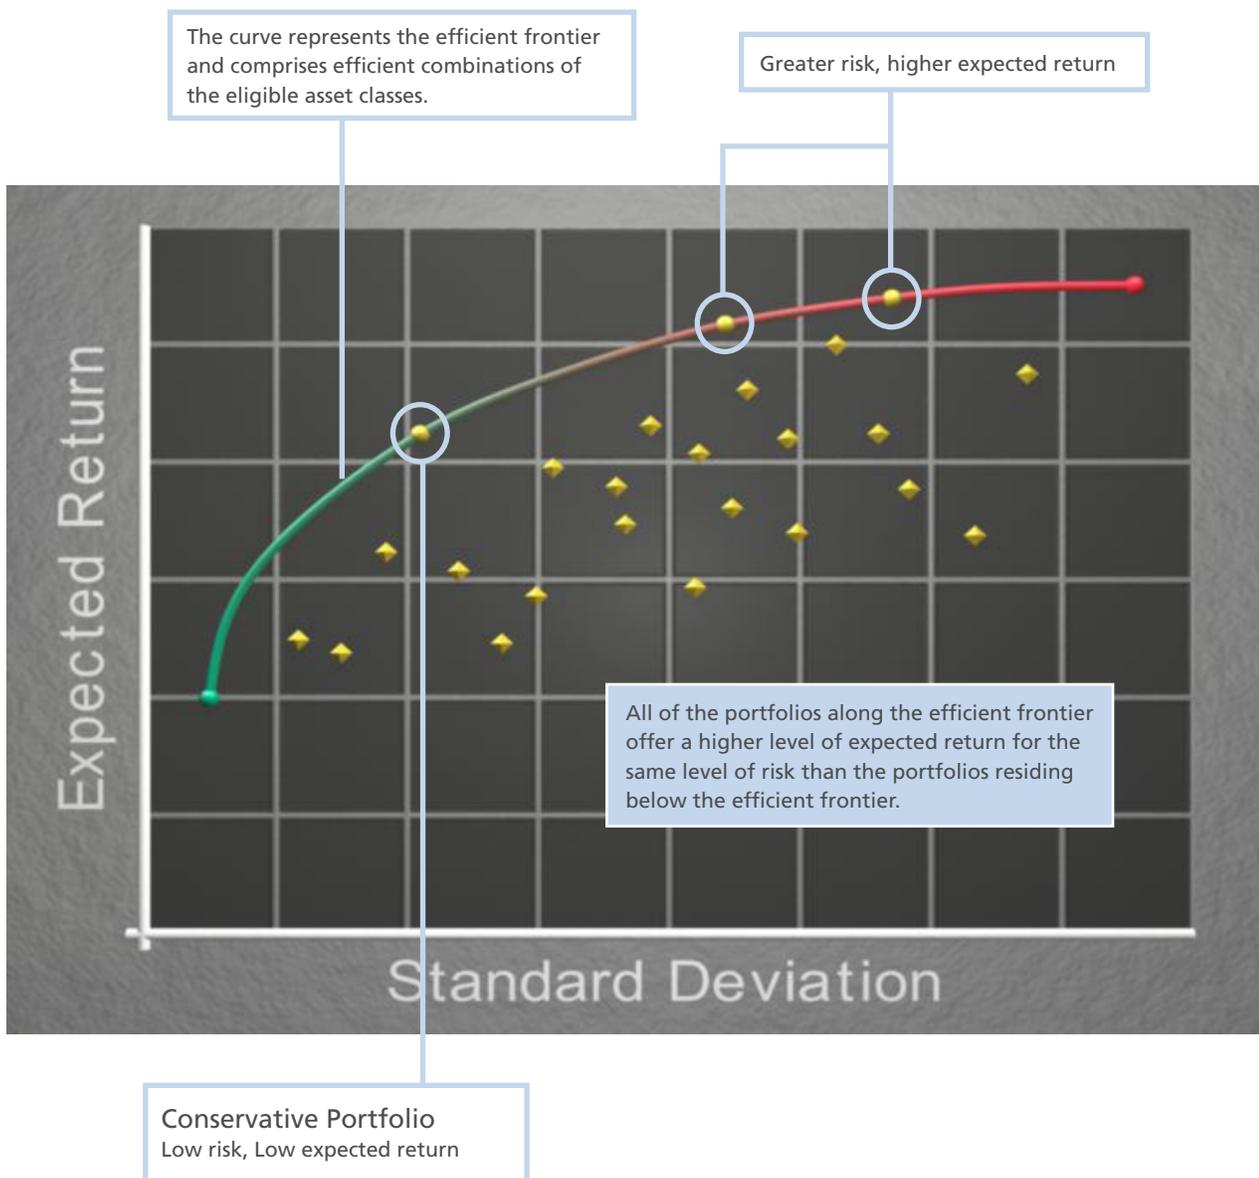
Normal	U.S. stocks	Foreign stocks	U.S. bonds	Real estate	Commodities
Foreign stocks	0.5704				
U.S. bonds	0.2468	0.1286			
Real estate	0.4594	0.3080	0.2041		
Commodities	-0.0327	0.1493	-0.0046	-0.0633	
Cash equivalents	0.0038	-0.0445	0.1968	-0.0477	-0.1682

Turbulent	U.S. stocks	Foreign stocks	U.S. bonds	Real estate	Commodities
Foreign stocks	0.5797				
U.S. bonds	0.1762	0.1220			
Real estate	0.5491	0.3379	0.2104		
Commodities	-0.0590	0.0229	0.0665	-0.1007	
Cash equivalents	0.0580	-0.0190	0.1170	0.0627	-0.1307

03 Efficient Portfolios

With this information, we use optimization to combine asset classes efficiently, so that for a particular level of expected return the efficiently combined asset classes offer the lowest level of risk, measured as standard deviation. A continuum of these portfolios plotted in dimensions of expected return and standard deviation is called the efficient frontier. Figure 2 shows a hypothetical efficient frontier.

Figure 2: Efficient frontier



Based on our earlier assumptions for expected returns, standard deviations, and correlations, we have derived three specific efficient portfolios: one for a conservative investor, one for an investor with a moderate appetite for risk, and one for an aggressive investor. We used the standard deviations and correlations from the normal periods to derive these portfolios. We use the risk values from both the normal and turbulent periods to measure their exposure to loss. Table 4 shows the three optimal portfolios.

Table 4: Optimal Portfolio

Asset Classes	Conservative	Moderate	Aggressive
U.S. stocks	22.86%	35.23%	48.15%
Foreign stocks	16.59%	24.22%	32.19%
U.S. bonds	49.95%	32.81%	14.89%
Real estate	3.85%	2.59%	1.28%
Commodities	6.75%	5.16%	3.49%
Cash equivalents	0.00%	0.00%	0.00%
Return	7.60%	8.37%	9.17%
Risk	7.77%	10.12%	12.86%

The composition of these portfolios should not be surprising. The conservative portfolio has nearly a 50% allocation to U.S. bonds. The moderate portfolio is well diversified and not unlike many institutionally-managed portfolios. The aggressive portfolio, by contrast, has more than an 80% allocation to U.S. and foreign stocks. It should be comforting to note that we did not impose any constraints on the optimizer to arrive at these portfolios. The process of employing an equilibrium perspective for estimating expected returns yielded nicely behaved results.

04 The Optimal Portfolio

The final step is to select the portfolio that best suits our tolerance for risk, which we call the optimal portfolio. The theoretical approach for identifying the optimal portfolio is to specify how many units of expected return we are willing to give up to reduce our portfolio's risk by one unit. If, for example, we are willing to give up $\frac{1}{2}$ unit of expected return to lower portfolio variance (the squared value of standard deviation) by one unit, our risk aversion would equal $\frac{1}{2}$. Risk aversion is the reciprocal of risk tolerance. We would then draw a line with a slope of $\frac{1}{2}$ and find the point of tangency between this line and the efficient frontier (with risk defined as variance rather than standard deviation). The portfolio located at this point of tangency is theoretically optimal because its risk/return tradeoff matches our preference for balancing risk and return. In practice, however, we do not know intuitively how many units of return we are willing to sacrifice in order to lower variance by one unit. Therefore, we need to translate combinations of expected return and risk into metrics that are more intuitive.

Because returns are approximately normally distributed, we can easily estimate the probability that a portfolio with a particular expected return and standard deviation will experience a certain loss over a particular horizon.⁴ Alternatively, we can estimate the largest loss a portfolio might experience given a certain level of confidence. We call this measure value at risk. We can also rely on the assumption of normality to estimate the likelihood that a portfolio will grow to a particular value at some future date.

Investors typically measure exposure to loss at the end of their investment horizon. This view of risk ignores what may happen along the way. Investors should think about risk differently. They should care about exposure to loss throughout their investment horizon and not just at its conclusion. We therefore focus on two additional risk measures to evaluate these portfolios: within-horizon probability of loss and continuous value at risk.⁵

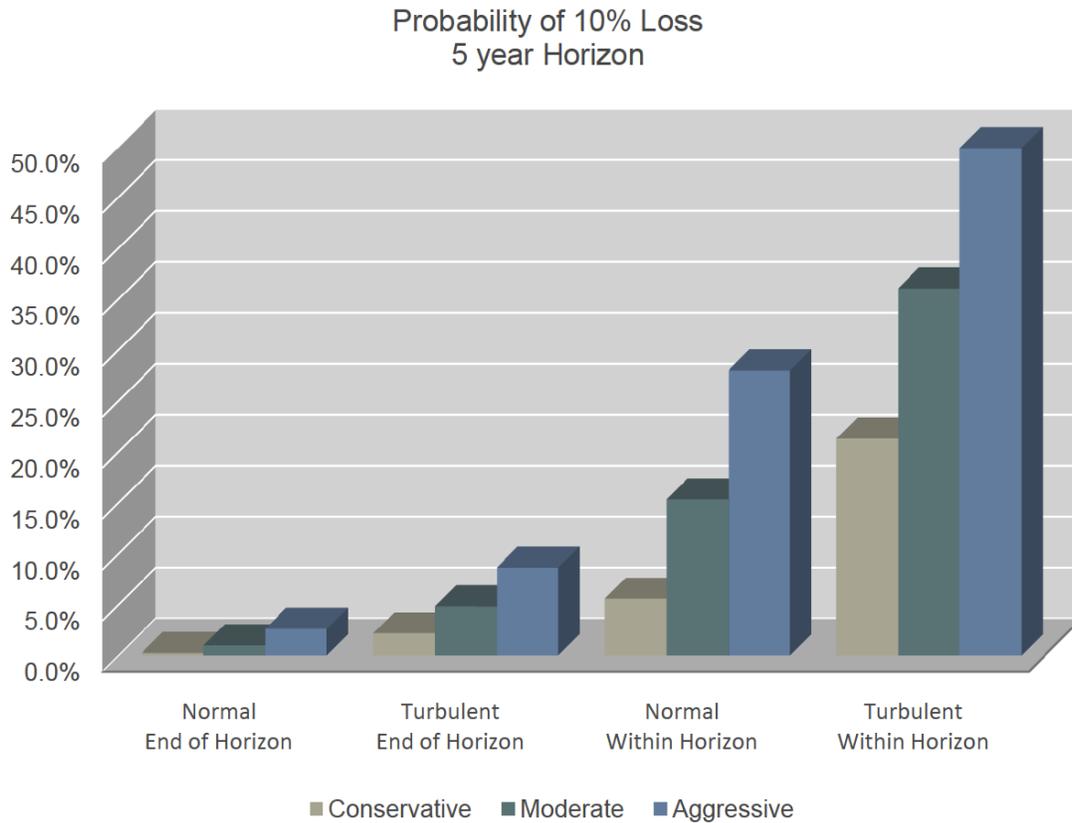
Within-horizon probability of loss measures the likelihood that an investment will depreciate to a particular level from inception to any point during a specified horizon and not just at the conclusion. Value at risk measured conventionally gives the worst outcome at a chosen probability at the end of an investment horizon. By contrast, continuous value at risk gives the worst outcome at a chosen probability from inception to any time throughout an investment horizon. As we shall soon see, these two risk measures reveal that within-horizon exposure to loss is substantially greater than investors typically assume.

⁴In theory, returns are lognormally distributed owing to the effect of compounding.

⁵For more detail about these risk measures, see Kritzman, M. and D. Rich, "The Mismeasurement of Risk", *Financial Analysts Journal*, May/June 2002

Figure 3 shows the likelihood of a 10% loss over a five-year investment horizon for the three efficient portfolios. These probability measures are presented for a normal regime, in which the risk parameters are based on the entire sample of returns, and for a turbulent regime, which uses the turbulent sub-sample of returns.

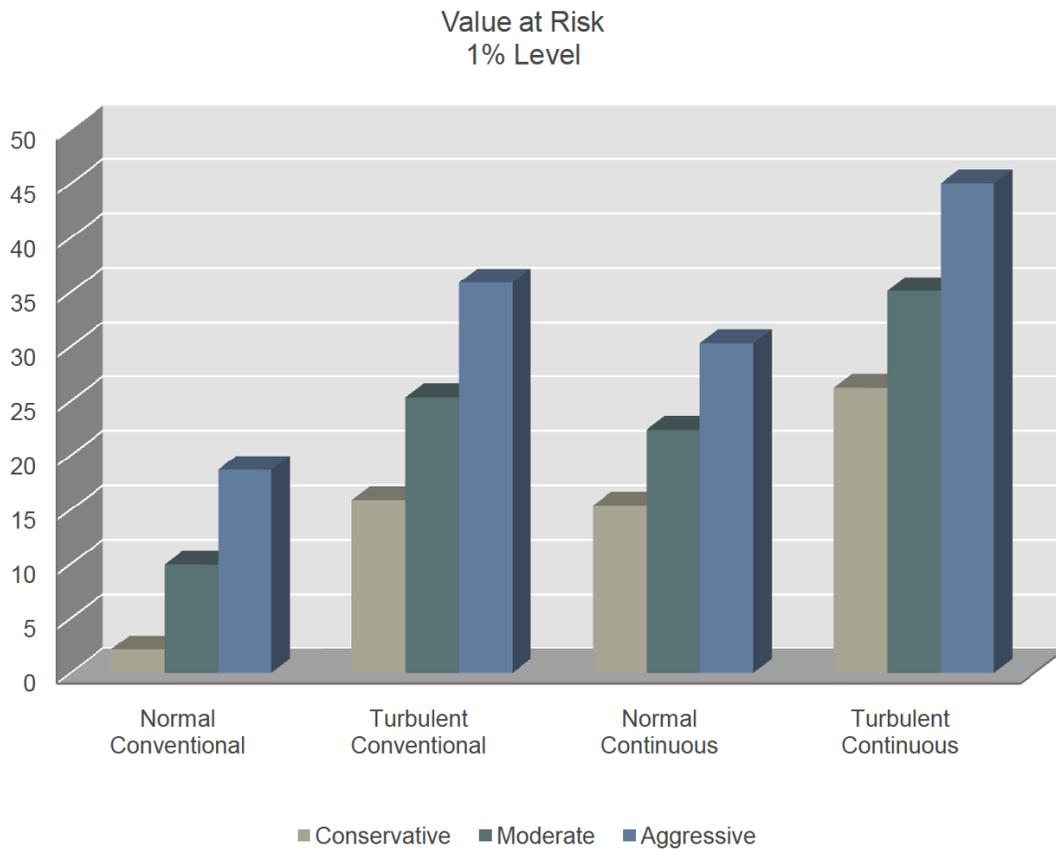
Figure 3: Probability of loss



End-of-horizon estimates of exposure to loss drastically understate a portfolio’s vulnerability to losses along the way. The moderate investor has only about a 1% chance of losing 10% or more at the end of five years, but there is a 15% chance that the portfolio will depreciate by at least that amount at some point along the way, and it increases to 36% if we expect a turbulent period to prevail. These are huge differences.

Figure 4 shows conventional value at risk (end-of-horizon) and continuous value at risk (within-horizon) for both a normal regime and a turbulent regime, measured at a 1% confidence level. Again, we see drastic difference in exposure to loss depending on whether we focus on the end of the horizon or the interim period as well. For example, the worst outcome at the end of the horizon for a moderate investor given a 1% probability and assuming a normal regime is a 10% loss. In comparison, the worst outcome at any point throughout the horizon is that the portfolio will decline by as much as 22%. If a turbulent regime prevails the worst outcome is a 35% loss.

Figure 4: Value at Risk



If exposure to loss were our only consideration, we would choose the conservative portfolio, but by doing so we might forego upside opportunity. One way to assess the upside potential of these portfolios is to simulate the distribution of future wealth associated with investment in each of them.

Suppose, for example, our current portfolio's value equals \$1,000,000 and we plan to contribute \$25,000 this year and grow our contributions annually by 5%. We can simulate how our portfolio's value will change through time with a procedure known as Monte Carlo simulation. This procedure generates possible future outcomes by drawing random numbers from a theoretical distribution such as a normal or lognormal distribution with a pre-specified expected return and standard deviation. If we wish to simulate outcomes that are relatively far into the future, it is important that we recognize the effect of compounding, which leads to a lognormal distribution rather than a normal distribution. Compared to a normal distribution, which is symmetrical, a lognormal distribution has a longer right tail than left tail and an average value that exceeds the median value.

To perform a Monte Carlo simulation of our wealth 15 years forward, assuming we invest in the moderate portfolio and contribute annually beginning with \$25,000 and growing this amount by 5% per year, we proceed as follows:

- ▶ We draw a random return from a lognormal distribution with an expected return of 8.37% and a standard deviation of 10.12%, the expected return and risk of the moderate portfolio.
- ▶ We then add 1 to the randomly selected return and multiply this value by \$1,000,000, the fund's initial value.
- ▶ Next we add the \$25,000 contribution.
- ▶ We then draw a second return from the theoretical distribution, add 1, and multiply this value by the new value of the fund including the contribution.
- ▶ Next we add our second contribution, which is now \$26,250.
- ▶ We repeat steps one through five until 15 random returns are selected and we record the ending value of the portfolio.
- ▶ We repeat steps one through six 5,000 times to generate 5,000 15-year sequences.
- ▶ Finally, we rank the 5,000 final values from highest to lowest and note the values at the percentiles in which we are interested.

Monte Carlo Simulation

Initial Wealth:

\$1,000,000

Investment Horizon:

15 Years

Annual Contribution:

\$25,000

**Annual Growth Rate
of Contributions:**

5%

Number of Simulations:

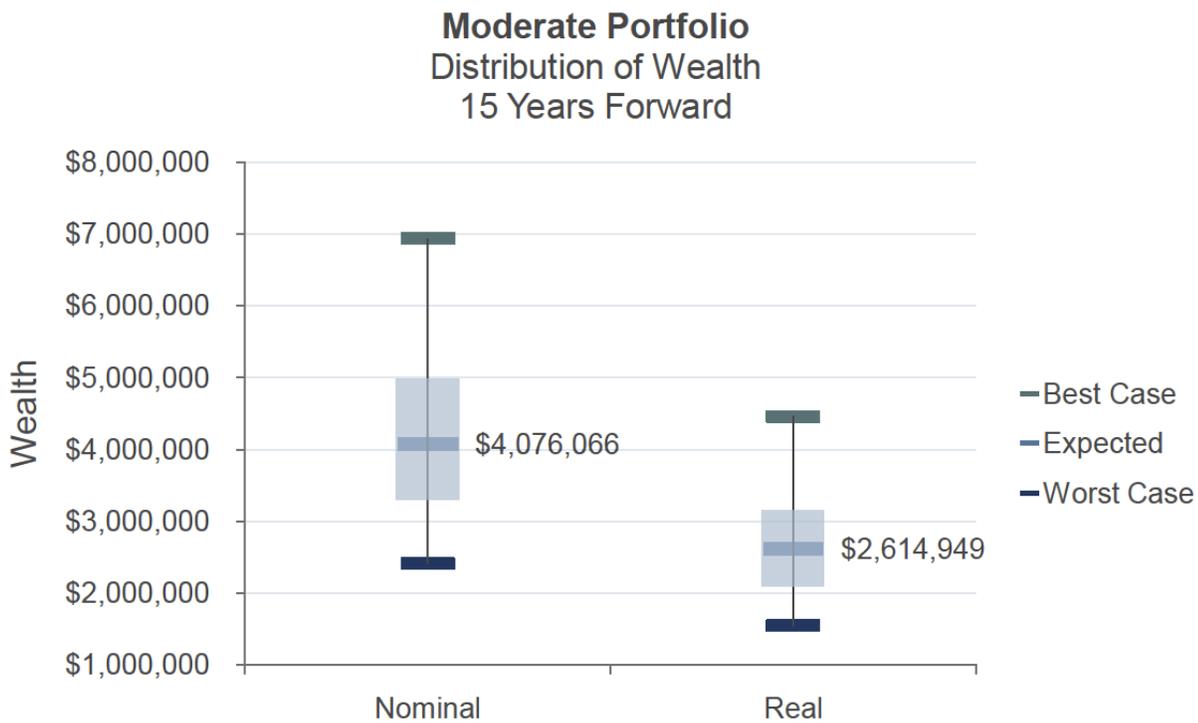
5000

Remainder at End Period:

\$250,000

Figure 5 shows the probable terminal wealth at different confidence levels in both nominal and real terms, assuming an inflation rate of 3.0%. For example, there is a 5% chance that our nominal wealth could grow to as much as \$7 million in fifteen years, assuming we invest in the moderate portfolio and make the annual contributions specified above. After inflation, this amount falls to about \$4.5 million. We should be 95% confident of ending up with about \$2.3 million which would be worth \$1.6 million after inflation. And we should expect to end up with \$4.1 million dollars nominally and \$2.7 after inflation.

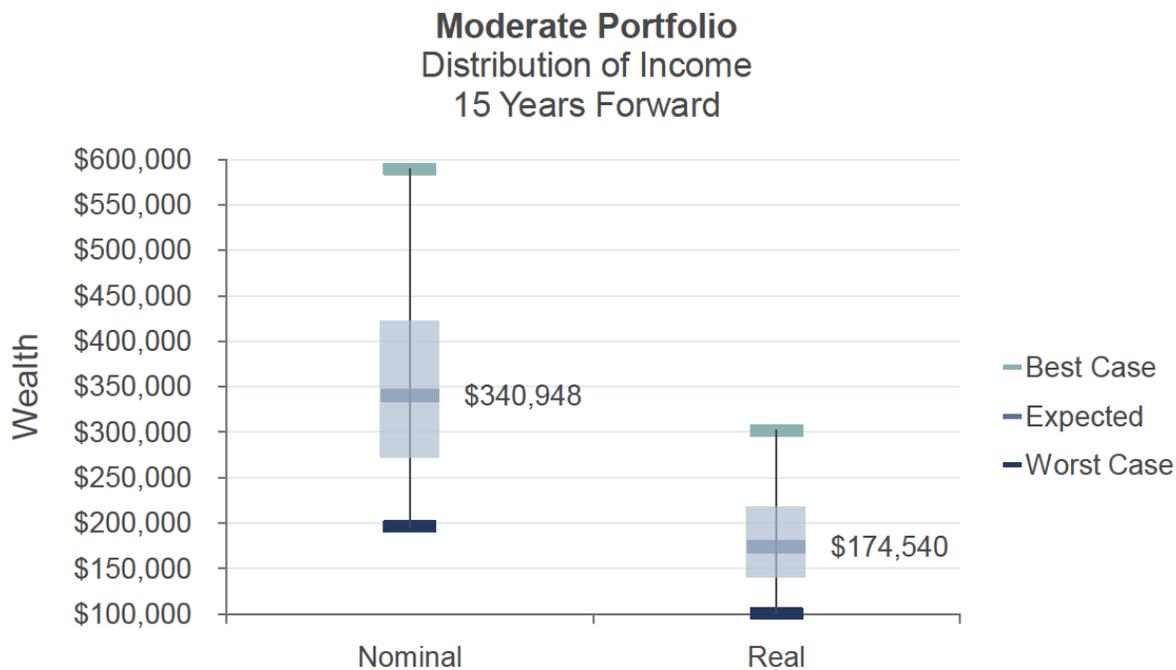
Figure 5: Distribution of Wealth



We can also use Monte Carlo simulation to estimate the distribution of our future income. We follow the same procedures as described above for estimating our future wealth, but we assume that the wealth at the end of each of the 5,000 paths is invested in an annuity at a particular interest rate for a specified period of time. In addition, we can specify an amount to be left at the conclusion of the annuity period.

Figure 6 shows the distribution of future income assuming an annuity rate of 3.5% and a remainder to be left of \$250,000. Our expected annual income in nominal terms is about \$340,000, and adjusted for inflation equals about \$175,000. There is a 5% chance it could be as high as \$590,000 and \$300,000 respectively, and we should be 95% confident that it will at least be \$195,000 nominally or \$100,000 adjusted for inflation.

Figure 6: Distribution of Income



By mapping the expected returns and standard deviations onto estimates of exposure to loss and the distribution of future wealth and income, we should have a clear idea of the merits and limitations of these portfolios. It is important to keep in mind, though, that there is no universally optimal portfolio; it is specific to each investor. If our focus is to avoid losses, the conservative portfolio might be optimal. If, instead, we believe that we can endure significant losses along the way in exchange for greater opportunity to grow wealth and future income, then we might choose the aggressive portfolio. If our goal is to limit exposure to loss, yet still maintain a reasonable opportunity to grow wealth and income, then perhaps the moderate portfolio would suit us best.

Asset allocation is a complex process, yet one we cannot ignore. We strongly urge investors to approach asset allocation with discipline, structure, and internal consistency. Moreover, we recommend that investors evaluate portfolio choices based not on statistical abstractions, but rather on intuitive interpretations of their risk and return attributes.