



Asset Allocation Methodologies

by Tom Coyne



EXECUTIVE SUMMARY

- Asset allocation is both a process and a collection of methodologies that are intended to help a decision-maker to achieve a set of investment objectives by dividing scarce resources between different alternatives.
- Theory assumes that asset allocations are made in the face of risk, where the full range of possible future outcomes and their associated probabilities are known. In the real world this is rarely the case, and decisions must be made in the face of uncertainty.
- The appropriate asset allocation methodology to use, in part, depends on an investor's belief in the efficacy of forecasting. Assuming you believe that forecasting accuracy beyond luck is possible, there remains an inescapable trade-off between a forecasting model's fidelity to historical data and its robustness to uncertainty. Confidence in prediction also increases when models based on different methodologies reach similar conclusions. In fact, averaging the results of these models has been shown to raise forecast accuracy.
- The traditional methodology for asset allocation problems is mean-variance optimization (MVO), which is an application of linear programming that seeks to maximize return for any given level of risk. However, MVO has many limitations, including high sensitivity to input estimation error and difficulty in handling realistic multiyear, multiobjective problems.
- Alternative techniques include equal weighting, risk budgeting, scenario-based approaches, and stochastic optimization. The choice of which to use fundamentally depends on your belief in the predictability of future levels of risk and return.
- Although they are improving, all quantitative approaches to asset allocation still suffer from various limitations. For that reason, relatively passive risk management approaches such as diversification and automatic rebalancing occasionally need to be complemented by active hedging measures, such as going to cash or buying options.

INTRODUCTION

Everyone has financial goals they want to achieve, whether it is accumulating a target amount of money before retirement, ensuring that a pension fund can provide promised incomes to retirees, or, in a different context, achieving an increase in corporate cash flow. Inevitably, we do not have unlimited resources available to achieve these goals. We often face not only financial constraints, but also shortages of information, time, and cognitive capacity. In many cases, we also face additional constraints on how we can employ available resources to achieve our goals (for example, limits to the maximum amount of funds that can be invested in one area, or the maximum acceptable probability of a result below some threshold).

Broadly, these are all asset allocation problems. We solve them every day using a variety of methodologies. Many of these are nonquantitative, such as dividing resources equally between options, using a rule of thumb that has worked in the past, or copying what others are doing. However, in cases where the stakes are high, the

allocation problem is complicated, and/or our choice has to be justified to others, we often employ quantitative methodologies to help us identify, understand, and explain the potential consequences of different decision options. This article considers a typical asset allocation problem: how to allocate one's financial assets across a range of investment options in order to achieve a long-term goal, subject to a set of constraints.

THE CORE CHALLENGE: DECISION MAKING UNDER UNCERTAINTY

All investment asset allocation methodologies start with two core assumptions. First, that a range of different scenarios could occur in the future. Second, that investment alternatives are available whose performance will vary depending on the scenario that eventually develops. A critical issue is the extent to which a decision-maker believes it is possible to accurately predict future outcomes. Traditional finance theory, which is widely used in the investment management industry, assumes that both the full range of possible out-

comes and their associated probabilities are known to the decision-maker. This is the classic problem of making decisions in the face of risk.

However, when you dig a bit deeper, you find that this approach is based on some questionable assumptions. The obvious question is: how can a decision-maker know the full range of possible future outcomes and their associated probabilities? One explanation is that they understand the workings of the process that produces future outcomes. In physical systems, and even in simple social systems, this may be true. But this is likely not to be the case when it comes to investment outcomes. Financial markets are complex adaptive systems, filled with positive feedback loops and nonlinear effects caused by the interaction of competing strategies (for example, value, momentum, and passive approaches) and underlying decisions made by people with imperfect information and limited cognitive capacities who are often pressed for time, affected by emotions, and subject to the influence of other people. An investor can never fully understand the way this system produces outcomes.

Even without such causal understanding, an investor could still believe that the range of possible future outcomes can be described mathematically, based on an analysis of past outcomes. For example, you could use historical data to construct a statistical distribution to describe the range of possible future outcomes, or devise a formula for projecting a time series into the future. The validity of both these approaches rests on two further assumptions. The first is that the historical data used to construct the distribution or time-series algorithm contain sufficient information to capture the full range of possible future outcomes. The second is that the unknown underlying process that generates the historical data will remain constant, or only change slowly over time. Over the past decade, we have seen repeated evidence that in financial markets these two assumptions are not true, for example in the meltdown of the Long Term Capital Management hedge fund in 1998, the crash of the technology stock bubble in 2001, and the worldwide financial market panic in 2008. In these cases, models based on historical data failed to identify the full range of possible outcomes, or to accurately assess the probability of the possible outcomes

“A goal without a plan is just a wish.” Antoine de Saint-Exupéry





2



Topic • Best Practice

they identified. People will live with the consequences of these failures for years.

This is not to say that skilled forecasters do not exist, however. They certainly do. Unfortunately, it is usually easier to identify them with the benefit of hindsight (which also helps to distinguish between skill and luck) than it is to pick them in advance.

This discussion leads to an important conclusion. In the real world, asset allocators must make decisions not in the face of *risk*, but rather under conditions of true *uncertainty*, in which neither the full range of possible future outcomes nor their associated probabilities are fully known in advance. This has two critical implications. First, there is an inescapable trade-off between any forecasting model's fidelity to historical data and its robustness to uncertainty. The more carefully a model is backtested and tightly calibrated to accurately reproduce *past* outcomes, the less likely it is to accurately predict the future behavior of a complex adaptive system. Second, confidence in a forecast increases only when models based on differing methodologies (for example, causal, statistical, time-series, and judgmental forecasts) reach similar conclusions, and/or when their individual forecasts are combined to reduce the impact of their individual errors. In short, decision-making under uncertainty is much harder than decision-making under risk.

Asset Allocation: A Simple Example

Let us now move on to a more concrete, yet still simple, example to illustrate some key issues that underlie the most common asset allocation methodology in use today. Our quantitative data and results are summarized in the following table:

| | Asset A | Asset B |
|---|---------|---------|
| Year 1 return | 1% | 3% |
| Year 2 return | 5% | 7% |
| Year 3 return | 9% | 20% |
| Year 4 return | 5% | -5% |
| Year 5 return | 1% | 8% |
| Sample arithmetic mean | 4.2% | 6.6% |
| Standard error of the mean | 1.5% | 4.1% |
| Sample geometric mean | 4.1% | 6.3% |
| Sample standard deviation | 3.3% | 9.1% |
| Covariance of A and B | 0.12% | |
| Correlation of A and B | 0.41 | |
| Asset weight | 40% | 60% |
| Expected arithmetic annual portfolio return | 5.6% | |
| Expected portfolio standard deviation | 6.1% | |
| Expected geometric annual portfolio return | 4.9% | |

Our portfolio comprises two assets, for which we have five years of historical data. In line with industry norms, we will treat each data point as an independent sample (i.e. we will assume that no momentum or mean-reversion processes are at work in our data series) drawn from a distribution which includes the full range of results that could be produced by the unknown return-generating process. As you can see, the sample mean (i.e. arithmetic average) annual return is 4.2% for Asset A and 6.6% for Asset B. So it is clear that Asset B should produce higher returns, right? Wrong. The next line of the table shows the standard error for our estimate of the mean. The standard error is equal to the sample standard deviation (which we'll discuss below) divided by the square root of the number of data points used in the estimate (in our case, there are five). Assuming that the data come from a normal distribution (that is, one in the shape of the bell curve), there is a 67% chance that the true mean will lie within plus or minus one standard error of our sample mean, and a 95% chance that it will lie within two standard errors. In our example, the short data history, along with the relatively high standard deviation of Asset B's returns, means that the standard errors are high relative to the sample means, and we really can't be completely sure that Asset A has a higher expected return than Asset B. In fact, we'd need a lot more data to increase our confidence about this conclusion. Assuming no change in the size of the standard deviations, the size of the standard error of the mean declines very slowly as the length of the historical data sample is increased—the square root of 5 is about 2.2; of 10, about 3.2; and of 20, about 4.5. Cutting the standard error in half—that is, doubling the accuracy of your estimate of the true mean—requires about a fourfold increase in the length of the data series. Considering that 20 years is about the limit of the available data series for many asset classes, you can see how this can create problems when it comes to generating asset allocation results in which you can have a high degree of confidence.

The next line in the table, the sample geometric mean, highlights another issue: As long as there is any variability in returns, the average return in a given year is not the same as the actual compound return that would be earned by an investor who held an asset for the full five years. In fact, the realized return—that is, the geometric mean—will be lower, and can quickly be approximated by subtracting twice the standard deviation squared from the arithmetic mean. In summary, the higher the vari-

ability of returns, the larger the gap will be between the arithmetic and the geometric mean.

The following line in the table shows the sample standard deviation of returns for Assets A and B. This measures the extent to which they are dispersed around the sample mean. In many asset allocation analyses, the standard deviation (also known as volatility) is used as a proxy for risk. Common sense tells you that the correspondence between standard deviation and most investors' understanding of risk is rough at best. Most investors find variability on the downside much less attractive than variability on the upside—and they like uncertainty even less than risk, which they can, or think that they can, measure. Also, when it comes to the distribution of returns, it is not just the average and standard deviation that are of interest to investors. Whether the distribution is Gaussian (normal)—that is, it has the typical bell curve shape—is also important. Distributions that are slightly tilted toward positive returns (as is the case with Assets A and B) are preferable to ones that are negatively skewed. Skewness should also affect preference for distributions with a higher percentage of extreme returns than the normal distribution (i.e. ones with high kurtosis). Preference for higher kurtosis should rise as skewness becomes more positive, and fall as it becomes more negative (i.e. as the probability of large negative returns rises). In fact, in our example, Asset B has positive skewness and higher than normal kurtosis (compared to Asset A's lower than normal kurtosis). Hence, some investors might be willing to trade off higher positive skewness and kurtosis against higher standard deviation in their assessment of the overall riskiness of Asset B. This might be particularly true when, as in the case of some hedge fund strategies, the expected returns on an investment have a distribution that is far from normal. However, many asset allocation methodologies still do not take these trade-offs into account, because they either assume that the returns on assets are normally distributed, or they assume that investors only have preferences concerning standard deviation, and not skewness or kurtosis.

Covariance and correlation

Covariance and correlation are two ways of measuring the relationship between the time series of returns on two or more assets. Covariance is found by multiplying each year's return for Asset A by the return for Asset B, calculating the average result, and subtracting from this the product of the average return for Asset A and by the

“By failing to prepare, you are preparing to fail.” Benjamin Franklin

QFINANCE





average return for Asset B—or, more pithily, it is the average of the products less the product of the averages. Correlation standardizes the covariance by dividing it by the product of the standard deviation of Asset A's returns, multiplied by the standard deviation of Asset B's returns. Correlation takes a value between minus one (for returns that move in exactly opposite directions) and plus one (for returns that move exactly together). In theory, a correlation close to zero implies no relationship between the returns on the two sets of returns. Unfortunately, most people forget that correlation only measures the strength of the *linear* relationship between variables; if this relationship is *nonlinear*, the correlation coefficient will also be deceptively close to zero. Finally, covariance and correlation measure the average relationship between two return series; however, their relationship under extreme conditions (i.e. in the tails of the two return distributions) may differ from this average. This was another lesson taught by the events of 2008.

Forming a Portfolio

Let us now combine Asset A and Asset B into a portfolio in which the first has a 40% weight and the second has a 60% weight. The second-to-last row of our table shows the expected arithmetic portfolio return of 5.6% per year. This is simply the weighted average of each asset's expected return. The calculation of the expected standard deviation of the portfolio is more complicated, but it highlights the mathematical logic of diversification. The portfolio standard deviation equals the square root of the portfolio variance. The latter is calculated as follows: [(Asset A weight squared multiplied by Asset A standard deviation squared) plus (Asset B weight squared multiplied by Asset B standard deviation squared) plus (two times Asset A weight multiplied by Asset B weight times the covariance of A and B)]. As you can see, the portfolio standard deviation is 6.1%, which is less than 6.8%—the weighted average of Asset A's and Asset B's standard deviations. The cause of this result is the relatively low covariance between A's returns and B's returns (or alternatively, their relatively low correlation of 0.41). The fact that their respective returns apparently move in less than perfect lockstep with each other reduces the overall expected variability of the portfolio return. However, this encouraging conclusion is subject to two critical caveats. First, it assumes the absence of a nonlinear relationship between A's returns and B's returns that has not been picked up by the correlation

estimate. Second, it assumes that the underlying factors giving rise to the correlation of 0.41 will remain unchanged in the future. In practice, however, this is not the case, and correlations tend to be unstable over time. For example, in 2008, investors discovered that despite relatively low estimated correlations between their historical returns, many asset classes shared a nonlinear exposure to a market liquidity risk factor. When liquidity fell sharply, correlations rose rapidly and undermined many of the expected benefits from portfolio diversification.

Expected Portfolio Returns

The last line in our table is an estimate of the geometric or compound average rate of return that an investor might be expected to actually realize on this portfolio over a multiyear period, assuming that we have accurately estimated the underlying means, standard deviations, and correlations and that they remain stable over time (all questionable assumptions, as we have noted). As you can see, it is less than the expected arithmetic annual return. Unfortunately, too many asset allocation analyses make the mistake of assuming that the arithmetic average return will be earned over time, rather than the geometric return. In the example we have used, for an initial investment of \$1,000,000 and a 20-year holding period, this difference in returns results in terminal wealth that is lower by \$370,358, or 12.5%, than the use of the arithmetic average would have led us to expect. This is not a trivial difference.

ASSET ALLOCATION: ADVANCED TECHNIQUES

The basic methodology we have just outlined can be used to calculate asset weights that maximize expected portfolio return for any given constraint on portfolio standard deviation (or other measure of risk, such as value-at-risk). Conversely, this approach can be used to minimize one or more portfolio risk measures for any given level of target portfolio return. These are all variants of the asset allocation methodology known as mean-variance optimization (MVO), which is an application of linear programming (for example, as found in the SOLVER function in an Excel spreadsheet). Although MVO is by far the most commonly used asset allocation methodology, it is, as we have shown, subject to many limitations.

Fortunately, there are techniques that can be used to overcome some, if not all, of the problems highlighted in our example. We will start with alternatives to the MVO methodology, and then look at alternative

means of managing errors in the estimation of future asset class returns, standard deviations, covariances, and other model inputs.

Alternative Approaches to Portfolio Construction

The simplest alternative to MVO is to allocate an equal amount of money to each investment option. Known as the $1/n$ approach, this has been shown to be surprisingly effective, particularly when asset classes are broadly defined to minimize correlations (for example, a single domestic equities asset class rather than three highly related ones, including small-, mid-, and large-cap equities). Fundamentally, equal weighting is based on the assumption that no asset allocation model inputs (i.e. returns, standard deviations, and correlations) can be accurately forecast in a complex adaptive system.

Another relatively simple asset allocation methodology starts from the premise that, at least in the past, different investment options perform relatively better under different economic scenarios or regimes. For example, domestic and foreign government bonds and gold have, in the past, performed relatively well during periods of high uncertainty (for example, the 1998 Russian debt crisis and the more recent subprime credit crisis). Similarly, history has shown that inflation-indexed bonds, commodities, and commercial property have performed relatively well when inflation is high, whereas equities deliver their best performance under more normal conditions. Different approaches can be used to translate these observations into actual asset allocations. For example, you could divide your funds between the three scenarios in line with your subjective forecast of the probability of each of them occurring over a specified time horizon, and then equally divide the money allocated to each scenario between the asset classes that perform best under it.

When it comes to more quantitative asset allocation methodologies, research has shown that—at least in the past—some variables have proven easier to predict and are more stable over time than others. Specifically, relative asset class riskiness (as measured by standard deviation) has been much more stable over time than relative asset class returns. A belief that relative riskiness will remain stable in the future leads to a second alternative to MVO: risk budgeting. This involves allocating different amounts of money to each investment option, with the goal of equalizing their contribution to total portfolio risk, which can be defined using either



“Prediction is very difficult, especially about the future.” Niels Bohr





4



Topic • Best Practice

standard deviation or one or more downside risk measures (for example, draw-down, shortfall, semi-standard deviation). However, as was demonstrated by the ineffective performance of many banks' value-at-risk models during 2008, the effectiveness of risk budgeting depends on the accuracy underlying assumptions it uses. For example, rapidly changing correlations and volatility, along with illiquid markets, can and did result in actual risk positions that were very different from those originally budgeted.

The most sophisticated approaches to complicated multiyear asset allocation problems use more advanced methodologies. For example, rather than a one-period MVO model, multiperiod regime-switching models can be used to replicate the way real economies and financial markets can shift between periods of inflation, deflation, and normal growth (or, alternatively, high and low volatility). These models typically incorporate different asset return, standard deviation, and correlation assumptions under each regime. However, they are also subject to estimation errors not only in the assumptions used in each regime, but also in the assumptions made about regime continuation and transition probabilities, for which historical data and theoretical models are quite limited.

Rebalancing Strategies

Multiperiod asset allocation models can also incorporate a range of different rebalancing strategies that manage risk by adjusting asset weights over time (for example, based on annual rebalancing, or maximum allowable deviations from target weights). When it comes to identifying the best asset allocation solution for a given problem, these models typically incorporate sophisticated evolutionary search techniques. These start with a candidate solution (for example, an integrated asset allocation and rebalancing strategy), and then run repeated model simulations to assess the probability that they will achieve the investor's specified objectives. An evolutionary technique (for example, genetic algorithms or simulated annealing) is then used to identify another potential solution, and the process is repeated until a stopping point is reached (which is usually based on the failure to find a better solution after a certain number of candidates have been tested or a maximum time limit is reached). Strictly speaking, the best solutions found using evolutionary search techniques are not *optimal* (in the sense that the word is used in the MVO approach)—meaning a unique solution that is, subject to the limits of the methodology, believed to be better

than all other possible solutions. In the case of computationally hard problems, such as multiperiod, multiobjective asset allocation, it is not possible to exhaustively evaluate all possible solutions. Instead, much as for real life decision-makers, stochastic search models aim to find solutions that are robust—ones that have a high probability of achieving an investor's objectives under a wide range of possible future conditions.

ESTIMATING ASSET ALLOCATION INPUTS

A number of different techniques are also used to improve the estimates of future asset class returns, standard deviations, correlations, and other inputs that are used by various asset allocation methodologies. Of these variables, future returns are the hardest to predict. One approach to improving return forecasts is to use a model containing a small number of common factors to estimate future returns on a larger number of asset classes. In some models, these factors are economic and financial variables, such as the market/book ratio, industrial production, or the difference between long- and short-term interest rates. Perhaps the best known factor model is the CAPM (capital asset pricing model). This is based on the assumption that, in equilibrium, the return on an asset will be equal to the risk-free rate of interest, plus a risk premium that is proportional to the asset's riskiness relative to the overall market portfolio. Although they simplify the estimation of asset returns, factor models also have some limitations, including the need to accurately forecast the variables they use and their assumption

that markets are usually in a state of equilibrium.

The latter assumption lies at the heart of another approach to return estimation, known as the Black–Litterman (BL) model. Assuming that markets are in equilibrium enables one to use current asset class market capitalizations to infer expectations of future returns. BL then combines these with an investor's own subjective views (in a consistent manner) to arrive at a final return estimate. More broadly, BL is an example of a so-called shrinkage estimation technique, whereby more extreme estimates (for example, the highest and lowest expected returns) are shrunk toward a more central value (for example, the average return forecast across all asset classes, or BL's equilibrium market implied returns). At a still higher level, shrinkage is but one version of model averaging, which has been shown to increase forecast accuracy in multiple domains. An example of this could be return estimates that are based on the combination of historical data and the outputs from a forecasting model.

When it comes to improving estimates of standard deviation (volatility) and correlations, one finds similar techniques employed, including factor and shrinkage models. In addition, a number of traditional (for example, moving averages and exponential smoothing) and advanced (for example, GARCH and neural network models) time-series forecasting techniques have been used as investors search for better ways to forecast volatility, correlations, and more complicated relationships between the returns on different assets. Finally, copula functions have been

▶ MAKING IT HAPPEN

- Using broadly defined asset classes minimizes correlations and creates more robust solutions by reducing the sensitivity of results to deviations from assumptions about future asset class returns, which are the most difficult to forecast.
- Equal dollar weighting should be the default asset allocation, as it assumes that all prediction is impossible.
- However, there is considerable evidence that the relative riskiness of different asset classes is reasonably stable over time and therefore predictable. This makes it possible to move beyond equal weighting and to use risk budgeting. There is also evidence that different asset classes perform better under different economic conditions, such as high inflation or high uncertainty. This makes it possible to use scenario-based weighting.
- Techniques such as mean–variance optimization and stochastic search are more problematic, because they depend on the accurate prediction of future returns. Although new approaches can help to minimize estimation errors, they cannot eliminate them or change the human behavior that gives rise to bubbles and crashes. For that reason, all asset allocation approaches require not only good quantitative analysis, but also good judgment and continued risk monitoring, even after the initial asset allocation plan is implemented.

“A man who does not think and plan long ahead will find trouble right at his door.” Confucius





employed with varying degrees of success to model nonlinear dependencies between different return series.

CONCLUSION

In summary, although they are improving and becoming more robust to uncertainty than in the past, almost all quantitative approaches to asset allocation still suffer from various limitations. In a complex adaptive system this seems unavoidable, since their evolutionary processes make accurate forecasting extremely difficult using existing techniques. This argues strongly for averaging the outputs of different methodologies as the best way to make asset allocation decisions in the face of uncertainty. Moreover, these same evolutionary processes can sometimes give rise to substantial asset class over- or undervaluation that is outside the input assumptions used in the asset allocation process. Given this, relatively passive risk management approaches such as diversification and rebalancing occasionally need to be complemented with active hedging measures such as going to cash or buying options. The effective implementation of this process will require not only paying ongoing attention to asset class valuations, but also a shift in focus from external performance metrics to achieving the long-term portfolio return required to reach one's goals. When your objective is to outperform your peers or an external benchmark, it is tempting to stay too long in overvalued asset classes, as many investors painfully learned in 2001 and again in 2008.

►► MORE INFO

Books:

Asset Allocation:

Bernstein, William. *The Intelligent Asset Allocator: How to Build Your Portfolio to Maximize Returns and Minimize Risk*. New York: McGraw-Hill, 2001.

Darst, David M. *The Art of Asset Allocation: Principles and Investment Strategies for Any Market*. 2nd ed. New York: McGraw-Hill, 2008.

Fabozzi, Frank J., Petter N. Kolm, Dessislava A. Pachamanova, and Sergio M. Focardi. *Robust Portfolio Optimization and Management*. Hoboken, NJ: Wiley, 2007.

Ferri, Richard A. *All About Asset Allocation: The Easy Way to Get Started*. New York: McGraw-Hill, 2006.

Gibson, Roger C. *Asset Allocation: Balancing Financial Risk*. New York: McGraw-Hill, 2000.

Michaud, Richard O., and Robert O. Michaud. *Efficient Asset Management: A Practical Guide to Stock Portfolio Optimization and Asset Allocation*. 2nd ed. New York: Oxford University Press, 2008.

Swensen, David F. *Pioneering Portfolio Management: An Unconventional Approach to Institutional Investment*. New York: Free Press, 2009.

Forecasting:

Mlodinow, Leonard. *The Drunkard's Walk: How Randomness Rules Our Lives*. New York: Pantheon Books, 2008.

Osband, Kent. *Iceberg Risk: An Adventure in Portfolio Theory*. New York: Texere, 2002.

Rebonato, Riccardo. *Plight of the Fortune Tellers: Why We Need to Manage Financial Risk Differently*. Princeton, NJ: Princeton University Press, 2007.

Taleb, Nassim Nicholas. *The Black Swan: The Impact of the Highly Improbable*. New York: Random House, 2007.

Articles:

There are many academic papers on asset allocation and portfolio construction methodologies. The best single source is www.ssrn.com. SSRN is also a good source for papers on markets as complex adaptive systems by authors including Andrew Lo, Blake LeBaron, Cars H. Hommes, and J. Doyne Farmer.

Websites:

In addition to web-based tools based on mean-variance optimization, there are many vendors of more sophisticated asset allocation software. All of the following employ advanced techniques beyond simple MVO:

AlternativeSoft: www.alternativesoft.com

EnCorr: corporate.morningstar.com/ib/asp/subject.aspx?xmlfile=1221.xml

New Frontier Asset Allocation Suite: www.newfrontieradvisors.com

SmartFolio: www.smartfolio.com

Windham Financial Planner: www.windhamcapital.com

"It is a bad plan that admits of no modification." Publilius Syrus