Asset Allocation with Liabilities – Redefining Risk

Wayne Pushka

May 1999

An asset allocation model is developed enabling investors to include liabilities directly into an optimization. In the process, risk is redefined as the possibility and consequences of failing to achieve one's goals. This is an extension of Telser's criteria. New efficient surfaces are built which differ substantially from the traditional efficient frontier. The method enables practitioners to align their investment strategies with their true objectives. It also assists them in defining their objectives and developing a better understanding of the risks they face.

The purpose behind investing is to achieve one's financial goals. Practitioners recognize asset allocation as crucial in this pursuit and studies have shown that it is responsible for the majority of portfolio behavior[1]. Current methods are deficient, however, in that they do not *directly* link it to investor goals. While most practitioners consider their goals, they do not incorporate them into their asset allocation optimization.

This article presents a generalized solution to the asset allocation problem. It describes a multiple period model allowing practitioners to enter time dependent risk, return, and correlation estimates for assets and liabilities directly into an optimization routine.

The application of Modern Portfolio Theory (MPT) involves the construction of an efficient frontier by optimizing the return / risk characteristics of various assets. When generating the frontier, most practitioners define return as expected return, and risk as the standard deviation of the expected return's dispersion (mean-variance). This dispersion is assumed to be normally distributed with the mean equal to the expected return. The frontier is calculated for a single period, usually one year, and the investor selects their desired point along the frontier based on their tolerance for risk.[2]

This method is used by sponsors and consultants to construct institutional portfolios. However, the method is essentially deficient in that it does not directly link investor goals with the means of achieving those goals. What is the point of investing in the first place if it is not to achieve one's financial goals?

This disconnection between investments and goals is reflected in the way MPT perceives risk. By not accounting for goals, or attempting to capture this complex relationship with a simple multiplier,[2] it eliminates a key aspect of risk. It becomes simply the mean variance of returns. This is a definition of uncertainty not risk, but most practitioners use it as their definition of risk since they lack a better alternative.

The simplification of risk has several weaknesses. First, it does not consider the link between a risky event and the consequences of that event. The movement of an asset's value is not risk, rather risk is manifest by the *consequences* of that asset's movement vis-à-vis the investor's goals. As the saying goes "It isn't the fall that kills you, it's the sudden stop." While an attempt can be made to build the event / consequence link into the utility function, the practitioner is faced with the daunting task of defining a function for each investor.

Second, risk is treated as having no direction in its dispersion from the mean. Consequently, high returns are considered just as risky as low returns. Investors respond differently to investment out performance than under performance. If investment returns exceed our forecast we normally accept this as good fortune, whereas uncertainty is treated as risk only when returns fall below expectations.

Finally, there is an assumption of normality in return distributions. While returns tend to be normally distributed in the short term, compounding transforms these normal distributions into lognormal ones. The longer the time period and higher the average return, the more a portfolio will deviate from this normal assumption.

In order to overcome these failings, risk is redefined. *Risk is the possibility and consequences of failing to achieve one's goals*. Financially, we treat goals as a future liability where they are defined, not as accounting entries, but as requirements for cash. For example, an insurance company's future liabilities include all future business expenses and claims less any premium income. While this definition is analogous to Telser's criteria, we have extended it by setting multiple goals instead of just one.

Our new definition of risk addresses many of the drawbacks found in traditional methods. By including goals into the definition, a link is established between the portfolio's volatility and the consequences of that volatility. This also results in low returns contributing to risk and high returns diminishing risk. The assumption of normal return distributions is not implicitly addressed in our new definition, rather the mathematics are constructed to accommodate non-normal possibilities.

The consequences of redefining risk are dramatic. Prior to exploring these results, we will briefly explain the model used for including liabilities and redefining risk. Readers not interested in the underlying mathematics and the redefinition's impact on the objective function may skip the next two sections and move to the Consequences.

The Model

A good model should be intuitive and tractable, while at the same time capturing the critical factors and relationships needed to accurately solve the problem. Unfortunately, there is usually a tradeoff between these objectives. When designing this model, we looked at the problem from the perspective of the practicing investor, with all the uncertainties and inaccuracies they face. By its nature, investing involves predicting future events. While one will rarely get a prediction exactly correct, what is more important is the degree to which forecasting errors impact the investor's circumstances. It is with the future's uncertainty that the investor must struggle.

When applying this method, the practitioner should keep in mind the idea of frame of reference or perspective. As we exist only in the present, future probability distributions are also viewed from the present, not from a point in the future. Events in time period t_n are viewed from the start period (the present) t_0. As a result, the expected returns and variances for each time period are cumulative, where $t_n < t_{n+1}$ with each period starting at t_0. A non-cumulative view would be with each period t_{n+1} starting at t_n (see Figure 1). The cumulative view embeds the effect of compound interest into the statistics and builds a causal link directly into the optimization. (In the following discussion, the terms t_1, t_2, ... , t_n are defined in the non-cumulative view and converted into the cumulative view within the expected return function.)

Figure 1. Cumulative vs. Non-Cumulative View

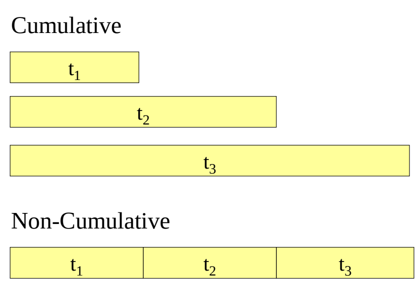

This depicts the investor's reality. When looking forward in time, decisions are made on what is known today, not on what will be known sometime in the future. This approach has two major side effects on the model, one negative and the other positive. The negative effect is that it greatly increases the difficulty in the initial mathematics. We are no longer dealing with normal (or near normal) distributions, but are now dealing with compounded (near) normal distributions. The positive effect is that the causal link results in the number of computations expanding linearly rather than geometrically, over time.

Including Liabilities

We begin by including our liability forecasts directly into our expected return. Rather than viewing future liabilities as negative investments (assets), we consider them to be negative returns on our existing assets. They are future obligations that are met with the future returns on existing assets (as well as cash from other sources such as dues, premiums or other income). Therefore, the total market value of our portfolio at each time period becomes the compounded return of our assets less the reductions due to our liabilities. The expected value of the portfolio then becomes:

$$E(T_1) = [1 + E(\alpha_1)]T_0 - E(L_1)T_0$$

$$E(T_2) = [1 + E(\alpha_2)]E(T_1) - E(L_2)T_0$$

…

$$E(T_t) = [1 + E(\alpha_t)]E(T_{t-1}) - E(L_t)T_0$$

where $E(T_t)$ is the expected portfolio value in period t; $E(\alpha_t)$ is the expected portfolio return for period t; $E(L_t)$ is the expected "return" on liabilities in period t; and T_0 is our portfolio's initial starting value. Since we know the portfolio's initial value, there is no expected value or uncertainty in the value of T_0. Also, since the portfolio value in time period t is dependent upon its value in t-1, a causal link is embedded directly into the algorithm. This is where the cumulative view enters the mathematics.

Expected returns on our assets can be further divided into returns and weights for individual assets, such that:

$$E(\alpha_t) = \omega_1 E(\alpha_{t1}) + \omega_2 E(\alpha_{t2}) + ... + \omega_n E(\alpha_{tn})$$

where ω_i is the weight associated with asset i with the no borrowing constraint $\sum_n \omega_i = 1$, and the no short selling constraint $\omega_i \geq 0$.

We treat liability returns differently. They are characterized using the following multiple factor model:

$$E(L_1) = E(\xi_{11})(1 + E(\varphi_{11})) + E(\xi_{12})(1 + E(\varphi_{12})) + .. + E(\xi_{1m})(1 + E(\varphi_{1m})) + E(\varepsilon_1)$$

$$E(L_2) = E(\xi_{21})(1 + E(\varphi_{11}))(1 + E(\varphi_{21})) + E(\xi_{22})(1 + E(\varphi_{12}))(1 + E(\varphi_{22})) + ..$$
$$.. + E(\xi_{2m})(1 + E(\varphi_{1m}))(1 + E(\varphi_{2m})) + E(\varepsilon_2)$$

$$E(L_t) = \sum_{l=1}^{m} [E(\xi_{tl}) \prod_{k=1}^{t} (1 + E(\varphi_{kl}))] + E(\varepsilon_t)$$

where $\xi_{t\ell}$ is the weight associated with liability factor ℓ at time t; $\varphi_{k\ell}$ is the rate of return on factor ℓ at time k; and ε_t is a residual term at time t. The factor weights $\xi_{t\ell}$, are assumed to be independent of the factor returnsφ_{kt}. The asset and liability factor returns are assumed to follow a Markov Process, where the asset returns at t are independent of their returns at t-1, and are only dependent upon the portfolio value at t-1. The liability factor returns at t are also assumed to be independent of the factor returns at t-1 as well as the asset returns at t-1. These independence (or orthogonality) relationships of the random variables are summarized as follows:

$$\{\alpha_{11},...,\alpha_{1n},\varphi_{11},...,\varphi_{1m}\}\perp...\perp\{\alpha_{t1},...,\alpha_{tn},\varphi_{t1},...,\varphi_{tm}\}\perp\{\xi_{11}\}\perp...\perp\{\xi_{tm}\}\perp\{\varepsilon_1\}\perp...\perp\{\varepsilon_t\}$$

and $\{\varphi_{t1}\}\perp...\perp\{\varphi_{tm}\}$

In order to simplify the problem, asset and liability factor returns in each time step are assumed to be normally distributed. This is not a restrictive assumption, as the total distribution still captures the effects of compound interest. Having said this, the restriction is not mandatory and the methodology can be fitted to real data.

When building the model, the user must forecast the asset class returns, liability factor returns, and the correlations between them. The weights should be calculated based on a linear regression of the forecasted liability returns and the liability factor returns. One difficulty is the proper selection of liability factors. Since they provide a link between the underlying behavior of the liability and the assets, a linear relationship must exist not only between the factor and the liability, but with at least one asset as well. In order to keep the underlying statistics stable, the factors themselves must be stable, independent of each other, measurable, meaningful and have a linear relationship with the liability returns. Examples of factors are inflation, real GDP growth, etcetera.

Combining all of these terms we get a generalized solution of the expected value of the portfolio at time t in terms of T_0.

$$E(T_t) = \left[[1 + E(A_t^0)] - \sum_{j=1}^{t} E(L_j)\left(1 + E(A_t^j)\right) \right] T_0 \qquad (1)$$

Notice the way the liabilities are treated. *Each liability return is separated from the initial portfolio and treated as the starting value of a new portfolio whose value is negative.* These liability portfolios compound at the same rate as the asset portfolio and are then subtracted back from the asset portfolio. Please see Appendix 1 for an example of how to use this method to determine the expected portfolio value. To find a multi-period solution, it is also necessary to define the intended re-balancing strategy and whether assets are perpetual (equities or a bond index) or mature (bond).

There is no consideration of the risks associated with each of the investments, nor the relationships between the assets and between assets and liabilities. In order to construct a complete solution, the concept of risk must be clearly defined.

Redefining Risk

Risk has been redefined as the possibility and consequences of failing to achieve one's goals. In order to accomplish this mathematically, minimum boundaries (or floors) are set for the portfolio. The floors are labeled F_{pt}, where p is the probability of the portfolio falling below the floor at each time period t. There is no limit to the number of floors that can be set for a given time period. Therefore, a ladder of floors can be used, each with a different probability. The floors, with their associated probabilities, act as tolerances for short and long term volatility, catastrophic events and as a means of setting the practitioner's confidence in their forecasting ability. Risk is now defined by multiple variables instead of just one.

There is now sufficient information to find a solution. To do this, a generalized density function for equation (1) is built. This function is the product of all the density and joint density functions of the random variables $\underline{\alpha}$, $\underline{\xi}$, $\underline{\varphi}$ and $\underline{\varepsilon}$. It includes a large variance-covariance matrix, which captures the relationships between the asset class returns and liability factor returns for each point in time.

The probability $P(T_t \leq F_{pt})$ is determined by solving the following integral:

$$\int \cdots \int_{\Re^{(n+2m+1)t}} f_{n,2m+1,t}(\underline{\alpha},\underline{\xi},\underline{\varphi},\underline{\varepsilon})\, d(\underline{\alpha},\underline{\xi},\underline{\varphi},\underline{\varepsilon}) \qquad (2)$$

where f is the generalized density function for T_t. The integral is fairly massive, with a dimension of (n+2m+1)t, where n is the number of asset classes, m the number of liability factors and t the number of time periods. So, for a problem with ten asset classes and two liability factors over a twenty year period (with time intervals of one year), a three hundred dimensional integral is built. The integral is solved numerically using a Monte Carlo method and the process repeated within an optimization routine.

The objective function is defined by maximizing expected return while satisfying the constraints. This function will be unique for each investor, since each investor has different constraints and expectations. The objective function itself is not known however, only its solution.

In practice, the investor must select their constraints and the corresponding probabilities. The constraints, or floors, is a concept that investors can intuitively understand. Selecting them isn't as difficult as it may appear and is more intuitive than deciding on the appropriate objective function in a utility equation. Nevertheless, when applying this method, the practitioner is faced with a rather daunting task. There is an enormous amount of initial overhead in terms of populating the various matrices and in maintaining the stability of the optimizations. The prudent practitioner still requires substantial experience in understanding the behavior of the underlying statistics and must gain an intuitive feel for the final results.

This method is particularly useful for exposing errors in the investor's understanding of the risks they face. Frequently, investors misperceive the effects of some risks. The model clearly identifies how various assumptions and risks impact the investor's true objectives. Some results are initially counterintuitive and by examining these results more closely, the investor can better appreciate their real risks and objectives.

Objective Function

Before discussing some consequences of this method, we will briefly review how it differs from traditional techniques. There are two general theories underpinning the creation of an objective function, one is the von Neumann Morgenstern (VM) theory and the other is classical utility theory.

VM theory is a powerful tool for decision making when faced with different possibilities[3]. Mathematically rigorous, it is one of the foundations of game theory. It involves the development of a characteristic function, which describes the game being played. This function is constructed by assigning a value for each 'play' (scenario) of the game. (Some authors have described these values as certainty equivalents.) For example, a single scenario game is a coin toss where 'heads' results in a $100 win and 'tails' in nothing. The value (or certainty equivalent) of this game is the maximum amount one would pay in order to play. Since the expected return is $50, most people would pay less than $50 to play the game (although if they were in a casino, they would pay more). A more complex game would involve multiple scenarios, with each assigned a value. A strategy is then a selection of scenarios. The optimal strategy has the combined scenarios with the highest value.

Utility theory approaches the problem differently. Instead of generating a characteristic function, a preference function is created for the investor. This is done by selecting all possible alternatives and ordering them by preference. Since we all prefer more wealth, a 'risk' term is included so that each outcome is a combination of return and risk. This results in an indifference curve that specifies the trade-offs an investor is willing to make between expected return and risk. Ultimately, the preference function and the VM characteristic function are the same, since they both describe the underlying tradeoff (objective function).

Methods based on these two theories have a number of serious drawbacks when building and managing real portfolios. First, they require decisions on the part of the investor that are not normally considered. In a simple coin toss it is easy to calculate the odds, but when making an investment decision no one calculates every possible scenario or ranks their preference and probability for every outcome. By the time an investor has considered every possibility, the situation has usually changed.

In order to overcome this drawback, the practitioner must make a number of approximations[4]. Problems arise when these approximations are inaccurate, or used inappropriately. Since both theories require the *explicit* creation of a function to find a solution, the practitioner typically will make a number of simplifications in order to build this function.

The multiple scenario lying at the core of both theories also creates computational problems. Utility theory requires multiple preferences and VM theory multiple "plays". The drawback to multiple scenarios is that the number of calculations required to find a solution increases geometrically over time. If one were to start with 100 scenarios for a single time period, after five time periods the number of possible scenarios would have expanded to 100^5 (10 billion). These expansions are unwieldy since they rapidly exhaust any computational power available. Once again, simplifications are possible, with their resulting loss of accuracy.

The solution in this article approaches the problem from a different perspective. Instead of explicitly creating an objective function, a unique function is *inferred*, based on the investor's constraints. The utility function is defined as the maximization of expected return, with boundary conditions. *The objective function is not known, rather its solution is numerically calculated based on the boundary conditions*. This method differs from the VM and classical utility based methods in four ways:

1. Utility is maximized by maximizing returns within a set of constraints.
2. There are no extraneous decisions required by the investor (certainty equivalents and preferences between outcomes are not required).
3. The objective function is not predefined in the solution. Only the behaviors of each underlying asset, liability and the investor's constraints are known.
4. Causal links are embedded into the objective function calculation. Therefore, multiple scenarios are not used. This results in the number of computations increasing linearly with time as opposed to geometrically.

Consequences of Redefining Risk

This redefinition of risk has a significant impact on Modern Portfolio Theory (MPT) and how it is applied to the investment management process. The two key variables in any aspect of MPT are return and risk, just as the two main concerns for an investor are future returns and the associated risks. By redefining one of these fundamental variables, a Pandora's box of issues is opened. We will explore some of these in a series of examples.

To reiterate, for readers who have skipped the previous two sections, the redefinition of risk is accomplished mathematically by setting minimum boundaries (or floors) to the portfolio. These floors can be set at any point in time and have an associated probability. The floor values capture the concept of goals and the associated probability captures the consequences of not achieving those goals.

Example 1 - Risk and the Possibility of Failure

The two aspects to our definition of risk are the possibility and consequences of failing to meet a goal. In this example, we will demonstrate how the traditional efficient frontier method does not capture the possibility of failing to meet a goal.

Suppose one can select from two assets. Asset 1 has an expected return of 10% with a range of possible returns of ±5%, while Asset 2 has an expected return of 4% with a range of ±1%. Using the mean variance definition of risk, (the probability distributions in this case are rectangular with variances of 8.33 and 0.33 respectively), our efficient frontier with no correlation between the two assets is shown in Figure 2.

Figure 2. Efficient Frontier for Two Assets

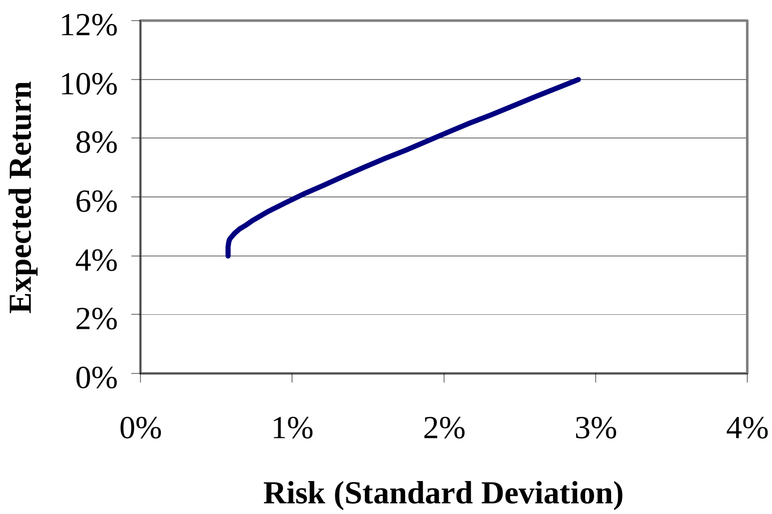

Using this efficient frontier, a naïve investor who is very risk averse would select 100% of Asset 2, since this asset has the lowest "risk". The problem arises in that dispersion of return is a poor proxy for measuring risk. The only "efficient" portfolio is 100% of Asset 1, since the worst case scenario for Asset 1 matches the best case scenario for Asset 2. Since our new definition of risk directly considers the probability distribution in the optimization, it selects 100% of Asset 1 as the only efficient portfolio.

In an efficient market, these two assets should not exist simultaneously, since no investor would purchase Asset 2 when Asset 1 will always out perform. The problem with the existing efficient frontier is that it offers Asset 2 as a possible solution. This is correct if we perceive risk as the variance of returns. However, uncertainty (or variance) is only one aspect of risk. Without explicitly including a goal when defining uncertainty, the de facto goal becomes minimizing the deviation from the expected return, or minimizing the uncertainty of the final outcome. But investors are not concerned with uncertainty about a mean, rather they are concerned with the uncertainty of meeting a goal.

In our example, the maximum possible return for Asset 1 is equivalent to the minimum possible return for Asset 2. Since our model considers the possibility of meeting a goal, it will always select Asset 1 as the optimal portfolio. Provided our investor is rational, no matter what the goal, Asset 1 has the highest likelihood of achieving it.

Example 2 - Risk's Two Dimensions

While the possibility of failing to meet a goal is one aspect of risk, the other is the consequences of that failure. With these two dimensions, we can no longer describe an efficient frontier in the traditional sense.

This is not a particularly revolutionary concept since many utility functions have a wealth factor in them that adds an extra dimension to risk. What is new, is the form and the way this extra dimension is added.

In this example, we describe our new efficient frontier using our new definition of risk. We begin by considering four assets with the expected returns and standard deviations listed in Table 1.

Table 1. Expected Returns and Standard Deviations of Four Assets

Asset	One Year Return	One Year Standard Deviation	Five Year Return	Five Year Standard Deviation
Equity 1	10%	14%	63.3%	31.3%
Equity 2	12%	17%	79.6%	38.0%
Bond	6%	7%	34.7%	15.7%
Money Market	4%	1%	22.1%	2.2%

Using five years as an investment horizon and initial portfolio value of one million dollars, the traditional method yields the efficient frontier shown in Figure 3.

Figure 3. Efficient Frontier: Four Assets

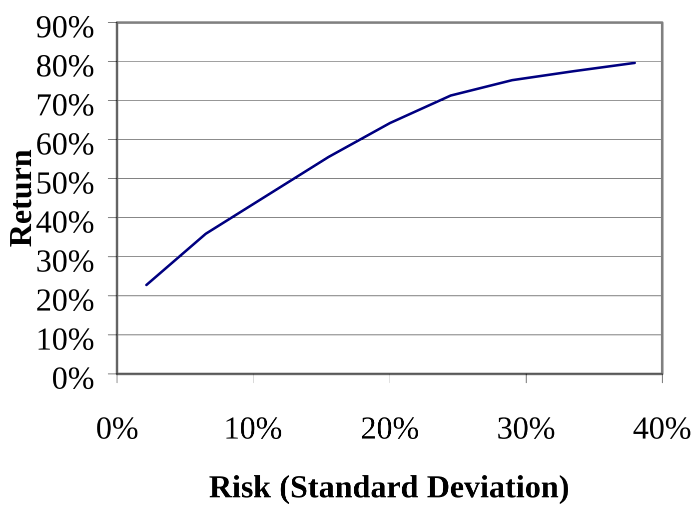

When applying our new method and definition of risk, we must have one more piece of information. We must know the investor's goal and their tolerance for failing to meet that goal. With this information, the method will then select the investor's optimal portfolio.

Most investors will require some help in converting their goal tolerance into a number. In order to assist them, we have constructed an efficient surface. By defining risk in terms of possibility and consequence, we have given it two dimensions. Adding expected returns results in a three dimensional plot. Figure 4 presents this new efficient frontier surface for our four assets.

Figure 4. Efficient Surface for Four Assets

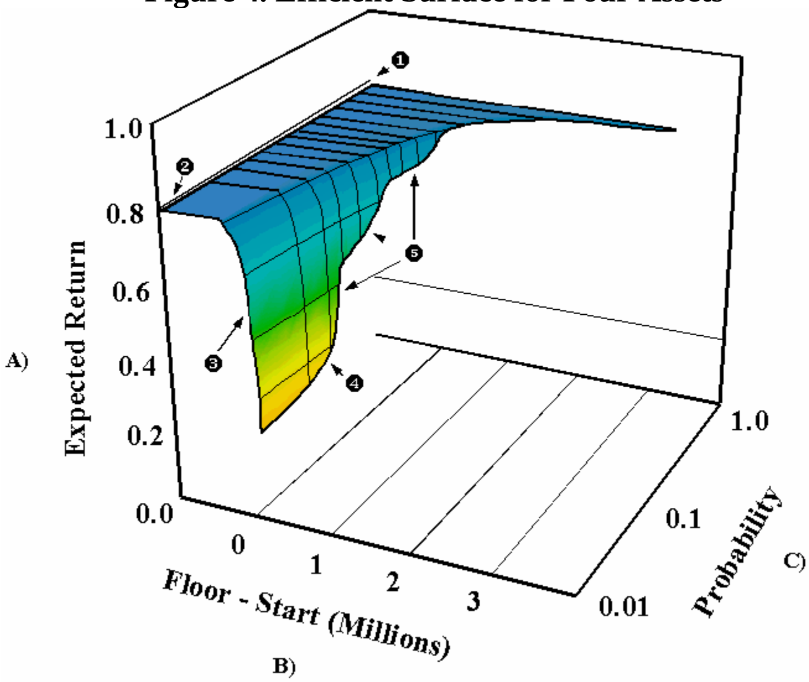

At first glance this surface plot looks rather intimidating. Three dimensional plots are difficult to grasp at first. Fortunately, most situations have very similar structures. Once the user becomes comfortable with interpreting this graph, they will have little difficulty in viewing other scenarios. The reader should avoid trying to equate the traditional efficient frontier with this new efficient surface. We have defined risk as the possibility and consequences of failing to meet a goal, while the traditional efficient frontier defines risk as the possibility of failing to meet the expected return.

Starting with the three axis:

A) Expected Return: The vertical axis is the same as in the traditional efficient frontier.

B) Floor - Start: This axis captures one component of risk. It is the difference between the floor value and the portfolio's starting value. The floor could be the investor's goal or a required distribution from a fund. In this problem we use only a single floor, where the floor is the final portfolio value. Since the axis is the difference between the two, if we have a start value of $1 million and a goal of $2 million, then the optimal portfolio lies above the $1 million axis ($2 million floor less $1 million start is $1 million). On the far left, the investor's portfolio is very large compared to their liability (or goal), and vice-versa on the far right.

C) Probability: This is the probability of falling below the floor. It is the investor's sensitivity of failing to meet their goal. The range is set from a 1% chance that the portfolio will fall below the floor, to a near certainty. So, an investor who sets a 99% probability for a goal is very risk tolerant and is speculating, while an investor with only a 1% probability has little tolerance for failure.

To help navigate the surface, we have highlighted the following features:

❶ The back left corner is the lowest risk position on the surface. The starting value of the portfolio is much higher than the floor and the investor's sensitivity to falling below the floor are minimal. The optimal portfolio is 100% of the highest returning asset.

❷ As one moves to the foreground, the consequences of falling below the floor become more severe. Even at a high sensitivity to falling below the floor (low probability), as long as the floor is substantially below the initial portfolio value (front left area of the graph) the optimal portfolio is still the highest returning asset.

❸ However, as the floor approaches and then exceeds the starting portfolio value, the likelihood of falling below the floor increases dramatically and the investor becomes more sensitive to the investments' volatility. There is a "waterfall" effect on the efficient surface as the asset classes with lower volatility and returns gain preference.

❹ As the floor is increased further (move right), the investor reaches a point where there is a very low likelihood that the low return asset can meet the higher target. Past this point (the "edge") there is no asset with a high enough return and low enough volatility to meet the higher floor target. Since there is no portfolio that can satisfy the investor's goals and risk tolerance, the investor must either lower their goals or increase their tolerance for risk. By lowering their goals, they will move to the left and by increasing their tolerance for risk, they will move back.

❺ There are "bumps" along the "edge". The inflection points between these "bumps" occur whenever another asset class is selected. The top area is 100% of the highest returning asset. The first "bump" is where two asset classes are selected, the second is where three are selected and the bottom "bump" contains all four asset classes.

Two dimensions to risk results in efficient surfaces as opposed to efficient curves. These surfaces can be very useful in understanding the interaction between risk and return.

Example 3 - Risk and the Consequences of Failure

So far we have considered the possibility of failing to meet a goal, and a means of describing the possibility and consequences of failure. This example will demonstrate why the consequences of failure must also be included in the definition of risk. In the process, it will also demonstrate how the lowest volatility is not necessarily the lowest risk.

Consider two assets, where Asset 1 has a low expected return with a high uncertainty of return ($E(R_1)$ = 10% and σ_1=50%), and Asset 2 has a high expected return with a low uncertainty of return ($E(R_2)$ = 15% and σ_2= 20%). In Figure 5, we plot this in the traditional way.

Figure 5. Risk vs. Return for Two Assets

High Return Low Volatility & Low Return High Volatility

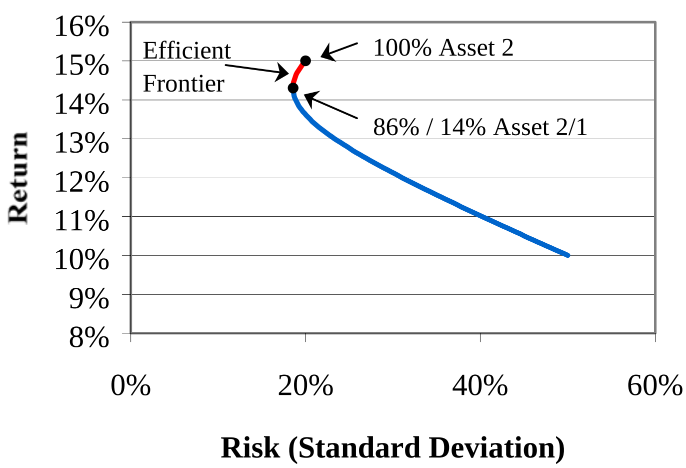

The "efficient frontier" in this situation is very short and is dominated by Asset 2. This results in a risk averse investor selecting a portfolio with 14% of the portfolio held in Asset 1 and 86% in Asset 2 and an aggressive investor opting for 100% invested in Asset 2. In the majority of situations this would be a rational decision. The situation changes if we view risk as the possibility and consequences of failing to meet one's goals.

In Figure 6, we have the difference between the initial portfolio value and the future liability on one axis, and the probability of falling below the liability on the other. This is an overhead view of our three dimensional plots. The three curves are Asset 1, Asset 2 and the minimum risk portfolio where we define risk as the likelihood and consequences of failing to meet a goal.

Figure 6. Risk Defined as Falling Below a Floor

High Return Low Volatility & Low Return High Volatility

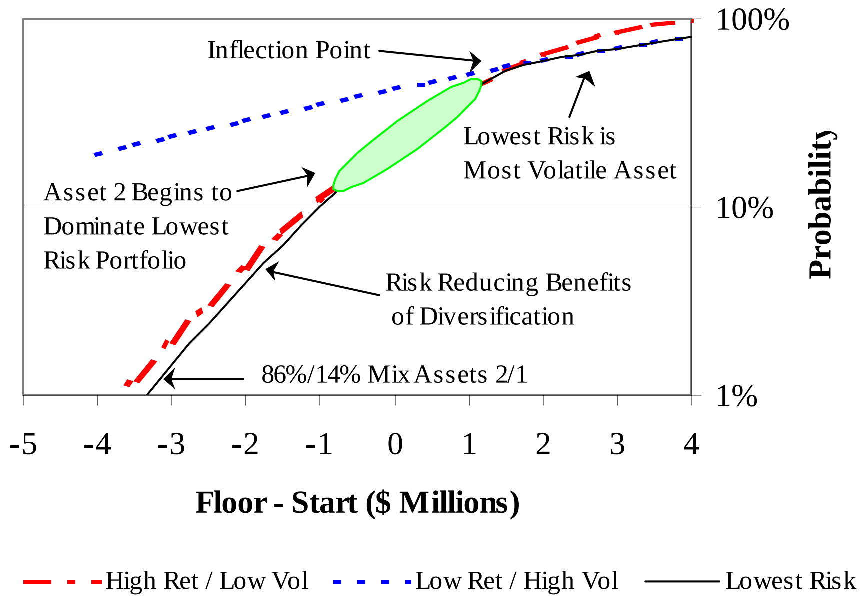

On the far left, the initial portfolio is substantially greater than the liability and the asset allocation decisions are the same as in the traditional approach; 14% in Asset 1 and 86% in Asset 2 for the risk averse investor, and a 100% position in Asset 2 for the aggressive investor. As the initial portfolio value approaches the liability's value (the center of the figure where the Floor - Start approaches zero), the portfolio becomes increasingly dominated by Asset 2. The risk reducing benefits of diversification are overcome by the low return and high volatility of Asset 1.

Interestingly, an inflection point occurs when the initial portfolio value falls to such a level that neither asset is likely to be sufficient in meeting the liability. To the right of the inflection point our investor becomes a desperado. They are more favorably disposed towards purchasing the riskier asset since that asset provides the highest likelihood of meeting the future liability. In this desperado situation, the rational investor prefers risk and the risk reduction benefits of diversification become a burden. If one owes a debt to a "loan shark" that is substantially greater than one's net worth, then it makes more sense to fly to Las Vegas and gamble in a casino with a negative expected return than to buy treasury bills. In the latter case one will be guaranteed to lose a limb, while the casino option at least provides the possibility of walking away unscathed.

In Example 2 we generated an efficient surface for four assets. If we add a fifth asset with a low return and high volatility, we get the surface in Figure 7. Notice the similarity to the surface in Figure 7 and the overhead view of two assets in Figure 6.

Figure 7. Efficient Surface for Five Assets with One Asset High Volatility Low Return

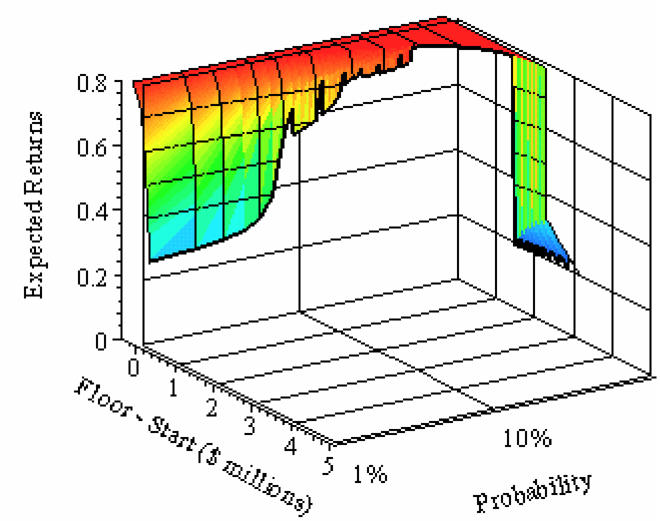

We can see the effect of the low return, high volatility asset on the right side of the surface. The inflection point is quite dramatic as this asset suddenly becomes the asset of choice. For example, suppose our desperado is starting with one million dollars and needs an additional four million. They would go to the four million point at the front of the graph and move back until they intersected a point on the efficient surface. In this case, this point would be 100% of the most volatile asset. Therefore, an asset with a high volatility could be the least risky asset because it offers the best chance of meeting the liability.

By including consequences in the definition of risk, what is initially perceived as an extreme gamble could in fact be risk averse behavior. Some behaviors that traditional theory would describe as irrational are now seen as being rational. By including consequences, we have a more realistic definition for risk and can more accurately describe investor behavior.

Example 4 - Risk and Complexity

Our redefinition of risk gives us the flexibility to solve a wide variety of more complex problems. For example, consider an individual who buys fire insurance on their home and at the same time purchases a few lotto tickets. By purchasing fire insurance, the individual is accepting the certain loss of a small sum (the premium) instead of the remote chance of a large loss (the value of the house). The expected return is negative, which is the cost incurred to buy certainty. Meanwhile, our individual has also purchased some lotto tickets, accepting the large chance of losing a small amount (the price of the tickets) in the remote chance of winning a large amount (the prize). The expected return is once again negative, but in this case it is the cost of purchasing uncertainty. So, on the same day, our individual is purchasing both certainty and uncertainty.

The original solution to this dilemma was the creation of a "wiggle" in the utility curve[5]. One difficulty with this solution is incorporating this "wiggle" into the standard efficient frontier model. By doing so, it adds a significant degree of complexity to the model. However, with our new method, we can effortlessly include this behavior.

Our individual's decision to purchase insurance and lotto tickets is based upon their need to meet future liabilities (goals). In this case there are two goals they must consider. The first goal is to maintain the value of their home. They want a 100% guarantee that the value of their home will not be destroyed by the unlikely event of fire. The home is necessary to fulfill the shelter component of their future living requirements. Since shelter is an important part of anyone's lifestyle, there is substantial suffering if it was to be destroyed. Our individual is prepared to pay for the certainty of meeting their future shelter requirements. An insurance instrument can be purchased (with a negative expected return) which will hedge against this event and therefore meet this constraint. At the other extreme is the lotto ticket. Our individual has another goal of living the rest of their life in comfort and luxury. They believe the only way they can achieve this goal is if they win the lotto. Since they do not perceive this goal as being "necessary" the consequence of failure is not high.

We have defined risk as the possibility and consequences of failing to achieve one's goals. Insurance is purchased as the lowest cost instrument that will guarantee the shelter goal is achieved, while the lotto ticket is perceived as the lowest cost instrument to achieve the goal of retiring in luxury. In both situations, an optimal instrument is being purchased to meet the constraints of the two goals. By combining these two goals and the two available investments, we find that the investor will purchase sufficient insurance to cover the cost of their home and purchase lotto tickets with the balance. If a third asset with positive expected return is added (a Treasury Bill for example), then our individual would hold a portfolio containing sufficient insurance and lotto tickets to meet their two constraints, with the balance in the positive returning asset.

Utility theory requires a complex function to solve this relatively simple problem. A real life situation would be even more difficult. By treating risk as a constraint, the model is more versatile and can easily accommodate these situations.

Example 5 - Risk and the Investment Horizon

This new definition of risk also has an impact on how the concepts of investment horizon and time diversification are viewed. Consider an investor who has $100,000 and wishes to purchase a house, requiring a down payment of $100,000. The investor can either invest these funds in a riskless asset such as Treasury bills, or a risky asset such as an S&P 500 index fund. One argument is that if an investor selects a riskless asset when planning to purchase a house in three months, then they should select a riskless asset when planning to purchase the house in ten years[6].

The problem with this argument is trying to determine what constitutes a "riskless" asset. Is a "riskless" asset a three month Treasury Bill, a ten year strip bond, or some other security? If we use one's ability to meet a goal as our definition of risk, then the S&P 500 index fund may be the better investment for our ten year horizon. When we define risk relative to a goal, in many cases there is no "risk free" rate of return or corresponding "riskless" asset. In this example, when a risk averse investor plans to purchase a house in ten years, they will seek a security that hedges their investment against changes in house prices. A ten year strip bond would provide a poor hedge against inflation while Treasury Bills are an imperfect hedge against asset inflation. If real estate experiences a burst of asset inflation over the ten year investment horizon, then the purchasing power of the investor's $100,000 would be greatly reduced. Therefore, it may be that of the three securities, the S&P 500 index fund provides the best hedge against asset inflation over the ten year investment horizon.

There are two sources of uncertainty facing the investor. One is uncertainty in the absolute value of the investment, and the other is uncertainty in the value of the liability (goal). The uncertainty in the value of a house in ten years time is quite significant, whereas (in most cases) it is likely to be fairly stable over three months. Therefore, if our investor's time horizon is only three months, then the three month Treasury Bill may be the best proxy for a "riskless" asset.

We can only find a risk free asset when there is a 100% certainty in the future value of the liability, or where the value of the liability and asset move in lock step with each other (perfectly correlated by direction and scale). These are formidable constraints, making risk free assets very difficult to find in most cases. We are left with trying to determine the asset that exhibits a minimum level of risk, rather than risk free.

This does not necessarily undermine the risk free rate in CAPM. That risk free rate is what the market views as being risk free, rather than any one individual investor. It is the weighted average minimum risk rate or return for all investors within that market. Problems arise in the asset allocation process when the investor assumes that this market risk free rate is equivalent to their own risk free rate.

Thus, risk and riskless assets are not a function of the absolute uncertainty in the investment's value. They are a function of the relative uncertainty of the investment. Since in many cases it is not possible to find a riskless asset, the investment horizon becomes relevant.

Example 6 - Relationships Between Assets and Liabilities

The success or failure of applying Modern Portfolio Theory methods depends upon how straightforward and robust these methods are when dealing with real world problems. Until now, we have only explored the consequences of our redefinition of risk using simple liability examples. Unfortunately, most liabilities are much more complex than this. Our method must be able to accommodate significantly more complex situations if it is to be useful.

Liabilities are treated differently than investments, since they have two characteristics. One is the timing and the amount of the liability and the other is any relationship that liability may have with the investments. The first characteristic involves estimating (with probabilities if necessary) the timing and amount of cash inflows and outflows. A future liability is some form of cash outflow (pension payments for a pension fund and insurance claims for an insurance company), and cash inflows (pension contributions for a pension fund) can be accommodated by simply treating them as negative liabilities. By capturing the second characteristic, the relationships between the liabilities and investments, the investor can build a superior portfolio.

In order to capture a relationship, it must first exist. The most obvious one is inflation, since nearly every liability and asset is effected by inflation. Others could be real GDP growth, or industry growth net of GDP and inflation. These relationships, or factors, must describe some part of a liability's behavior and part of at least one asset's behavior.

The problem in example two was repeated with an annual inflation rate of 2%, annual standard deviation of 1%, and a 90% correlation between the money market investment rate and inflation. The result, was a slight shift in the surface (since the liability became greater due to inflation), and a lengthening of the "waterfall" where the money market instrument was selected. Thus, there was a greater bias towards the money market instrument because of its high correlation with the inflation component of the liability.

An investor will want to find investments that have a positive correlation with their liabilities, since they act as natural hedges. By including liabilities into our definition of risk and considering the liability factor correlations with the investments, a "home bias"[7] will begin to arise.

Applications

This method was designed with the practitioner in mind. Instead of taking economic theory and applying it to investing, the model was constructed in an attempt to directly address the problems faced by the investor. By adopting this approach, the model is more robust and can accommodate real world "messy" data. It avoids extraneous or arbitrary inputs that can seriously distort the results, while using the same basic information to reach a decision as an investor would if they did not have this tool available. In this last section we touch upon a few situations where this method is applicable. They will be expanded upon in future articles.

Improving Investment Decision Making

By using this model the investor can develop a better understanding of the factors that drive their investing behavior. Whether the investor is an individual, corporation or institution, one of the main forces of investment behavior is the perception of risk. By developing a better understanding of that perception, the investor can design a superior strategy, thereby improving their long term performance.

One aspect of risk perception is the investment horizon. It is determined by when the investor needs their investment's proceeds. If they do not need the funds for a long period of time, then they should be relatively insensitive to short term fluctuations in their portfolio's value. In practice however, this often does not occur. There are several behaviors that suggest that investors frequently have a misalignment in their investment horizon.

Investors with significant money market (cash) components to their portfolios are a case in point. An organization with a high cash component is demonstrating either a high expected return for cash, or a short investment horizon. Frequently, investors with long investment horizons hold too much cash. Since liabilities are not directly included in the optimization, it is extremely difficult to ascertain how much is too much. A model such as this can help the investor see whether they have an inconsistency in their behavior towards their investment horizon.

Some behaviors do not have to be completely rational, but they must be recognized. For example, when an investor hires a new fund manager or embarks upon a new investment strategy, they have a very high aversion to seeing their portfolio's value fall below their initial investment. When their portfolio exceeds their initial investment, they become more risk tolerant. What is happening in this situation is that the investor perceives the initial investment value as a minimum value, ignoring the long term purpose of the funds and thereby ignoring the true investment horizon. This behavior is very common, and is irrational from an investment perspective. Nevertheless, it is reality and should be included in the design of an investment strategy. Value can be added by explicitly demonstrating to the investor the costs and consequences of this behavior and designing a strategy for minimizing its long term effects on the portfolio's value.

By explicitly addressing these behaviors and issues, discipline and structure is added to the investment decision making process. Too often decisions are made based on past prejudices and anecdotal evidence that do not stand up to careful scrutiny. Since there are so many factors and variables to consider (many with complex interrelationships) when making decisions, it is nearly impossible for the investor to sort through them all without some tools to aid them. Unfortunately, simple paradigms and beliefs are often the tools relied upon to make these decisions. By using a model that captures the relationships between investments and liabilities, and that defines risk more accurately, we can develop a much deeper understanding of the assumptions and perceptions driving an investor's decision making. By bringing objectivity to the process we ultimately improve portfolio performance.

Pension Fund Management

The management of pension funds is a natural fit for this method. Actuaries spend a great deal of effort determining future liabilities, thereby yielding significant amounts of information. The fund's objectives of providing pension benefits to retired employees are clearly defined, and the consequences of not meeting those objectives are also clear. With such specific information readily available, this method is well suited for determining an optimal pension portfolio.

The two main calculations for a pension fund are the ABO (Accumulated Benefit Obligation) and the PBO (Projected Benefit Obligation). The ABO calculation includes future cash flow estimates based on the obligations incurred from previous service, while the PBO calculation includes the ABO plus the expected benefits arising from future salary increases. Our model is an extremely powerful tool for determining the optimal portfolio under both the ABO and PBO assumptions. Instead of discounting the cash flows to a present value and using that as a target, the same underlying assumptions are used when forecasting both investment and liability growth. This improves the consistency of the decision making and generates a number of valuable synergies. These synergies are especially valuable when an investment shares the same underlying factors as a liability, thus providing a natural hedge.

The Private Investor

While the issues surrounding the investments for a private investor are similar to a pension fund, the liabilities can be much more ambiguous. The uncertainties surrounding the timing, consequences and magnitude of the future liabilities are frequently much greater in the case of the private investor.

The numerous uncertainties in life mean that future financial needs are very difficult to forecast. For example, in the case of an individual saving for retirement, the biggest uncertainty is their time of death. If a person dies on their 65th birthday they have no need for a retirement income, whereas living to 100 would require a substantial nest egg. Also, there is a tendency amongst people to constantly raise the bar of what they need. As one's wealth grows, an individual changes their needs from subsistence living to comfort and finally luxury. As they grow comfortable with a higher standard of living what were once perceived as luxuries are seen as necessities. This results in a raising of the financial goals as one's wealth increases.

At the extreme end, a wealthy investor's tolerance for risk should be quite high. According to traditional theory wealthy investors should have the majority of their investments in the highest returning assets. This is frequently not the case. In many instances, these investors become more loss averse, holding large portions of their portfolios in less volatile assets. If we assume rationality, then this behavior implies that while they have a higher likelihood of achieving their goals, their aversion to failure has also increased. This is an example of movement in both dimensions of risk.

One of the greatest challenges facing practitioners is assisting investors in determining their future goals. Setting those goals can be quite difficult, since most people have much greater latitude in deciding what to do with their wealth than do pension funds. They have a degree of flexibility and therefore uncertainty in terms of what their goals might be. Helping the client determine their goals, tolerance for risk and incorporating them into an investment strategy can be the most important service an investment adviser can give to their private client.

Dynamic Strategies

Dynamic strategies involve changing the asset mix of a portfolio when changes occur in an investor's risk profile or return forecasts. There are numerous strategies that can be adopted, including Constant-Mix, Portfolio Insurance and the less dynamic Buy-and-Hold[8]. These strategies are added to the traditional asset allocation method because the method is based on only a single time period, and therefore does not capture the effects of changes over time.

By using the technique described in this article, the investor can dispense with trying to select the proper dynamic strategy. Since this method captures the effects of liabilities over multiple time periods, we find that it incorporates the concept of dynamic strategies within its optimization. It is a generalized dynamic model where the various strategies of Constant-Mix, Portfolio Insurance, etcetera are all just specific solutions to a broader strategy.

When building a dynamic strategy, a series of scenarios should be run to determine how the reallocation would take place under different circumstances. Transaction costs can be included and an optimal dynamic strategy developed. By testing the strategy under various scenarios the investor can refine their initial conditions and underlying assumptions, improving their understanding of what drives their perception of risk.

Extensions

The model can be extended in a number of directions. In this article, we have only used single compounding rates of return. It is difficult enough to forecast a single long term return for an asset without trying to forecast returns for each time period. However, investors frequently have a short term expected return and a long term forecast. It is possible to adjust this method to account for this type of forecasting. Another issue, is that the model treats bonds as perpetual assets. A useful extension is to treat fixed term assets explicitly. The asset would then have a limited life and an uncertainty profile very different from a perpetual asset. This will require structuring a set of assumptions for the future reinvestment rate and incorporating it into the model. Applications could then be developed to assist insurance companies and banks in building strategies to manage their portfolios.

Conclusions

We have included liabilities into the asset allocation optimization and have redefined risk as the possibility and consequences of failing to achieve one's goals. By addressing many of the shortcomings of existing techniques, this method provides us with a powerful tool for building and managing investment portfolios. By focusing on goals, we can now design investment strategies oriented towards an investor's true objectives.

Notes

1. See Brinson, Hood & Beebower.

2. The resulting formula is the expected utility function. It has the form:

$$U = E(R) - \frac{\sigma_R^2}{\lambda}$$

   where E(R) is the expected return; σ_R^2 is the variance; and λ is a risk aversion multiplier.

3. For a detailed explanation of the method see the text by von Neumann and Morgenstern.

4. Ziemba and Mulvey have compiled a series of articles on solutions to Asset Liability Management problems. They have divided them into four methodologies; Stochastic Programming, Decision Rules, Capital Growth and Stochastic Control.

5. See the article by Friedman and Savage.

6. A 1994 Financial Analysts Journal article by Mark Kritzman puts forth this example.

7. Kang and Stulz discuss the home bias in international investing with Japan as a case in point. The basic assumption is if investors are only concerned with the mean and variance of the real return of their wealth, and with low barriers to international investing, then investors should hold the world market portfolio. This assumption assumes risk to be the variance of returns. Under our new definition of risk, there would be some bias towards holding a greater proportion of domestic assets.

8. Perold and Sharpe give an excellent discussion on these three basic dynamic asset allocation strategies.

References

Bierman, Harold J. 1997. “Portfolio Allocation and the Investment Horizon*.*” *Journal of Portfolio Management*, vol. 23, no. 4 (Summer): 51-55.

Bookstaber, Richard, and Gold, Jeremy. 1988. “In Search of the Liability Asset.” *Financial Analysts Journal*, vol. 44, no. 1 (January / February):70-80.

Brinson, Gary P., Hood, Randolph L., and Beebower, Gilbert L. 1986. “Determinants of Portfolio Performance.” *Financial Analysts Journal*, vol. 42, no. 4 (July/August): 39-44.

Elton, Edwin J., and Gruber, Martin J. 1991. *Modern Portfolio Theory and Investment Analysis 4 ed.* John Wiley & Sons, Inc.

Friedman, Milton, and Savage, Leonard J. 1948. “The Utility Analysis of Choices Involving Risk.” *Journal of Political Economy*, vol. 56, no. 4 (August): 279-304.

Hull, John C. *Options, Futures, and Other Derivative Securities*. 1993. Prentice-Hall International Inc.

Kang, Jun-Koo, and Stulz, Rene M. 1997. “Why is there home bias? An analysis of foreign portfolio ownership in Japan.” *Journal of Financial Economics*, vol. 46: 3-28.

Kritzman, Mark. 1994. “What Practitioners Need to Know … About Time Diversification.” *Financial Analysts Journal*, vol. 50, no. 1 (January / February):14-18.

Leibowitz, Martin L. 1986. “Total Portfolio Duration: A New Perspective on Asset Allocation.” *Financial Analysts Journal*, vol. 42, no. 5 (September / October):18-29, 77.

Leibowitz, Martin L., and Kogelman, Stanley. 1991. “Asset allocation under shortfall constraints.” *Journal of Portfolio Management*, vol. 17, no. 2 (Winter): 18-23.

Maginn, John L., and Tuttle, Donald L. 1990. *Managing Investment Portfolios: A Dynamic Process 2nd Ed*. Warren, Gorham & Lamont, Inc.

Markowitz, Harry. 1952. “Portfolio Selection.” *Journal of Finance*, vol. 7, no. 1 (March): 77-91.

Perold, Andre F., and Sharpe, William F. 1988. “Dynamic Strategies for Asset Allocation.” *Financial Analysts Journal*, vol. 44, no. 1 (January / February): 16-27.

von Neumann, J. and O. Morgenstern. 1944. *Theory of Games and Economic Behavior*. Princeton University Press.

Ziemba, William T. And Mulvey, John M. 1998. *Worldwide Asset and Liability Modeling*. Cambridge University Press.

Appendix 1

Consider the following scenario. An investor has $1000 to invest in three asset classes 1,2 and 3 with expected returns of 5%, 8% and 12% respectively. Assume the returns are compounded evenly over three years. At the end of year two the investor must pay $200 in current dollars and at the end of three years, $500 in current dollars. Assume the expected inflation rate is 3% per year.

Since the returns are compounded evenly, the expected returns on the assets are the same for each year. Therefore, $E(\alpha_1) = E(\alpha_2) = E(\alpha_3)$.

We have only assumed one liability factor (inflation) and there is no residual term. The factor is assumed to have the same expected compound rate for each year. Therefore, our liability factor rates of return are all equal; $E(\varphi_{11}) = E(\varphi_{21}) = E(\varphi_{31})$.

Our liability factor weights are calculated as follows;

for year one; 0
for year two; $200/$1000 = 0.2
for year three; $500/$1000 = 0.5

Table 2. Summary of Inputs

Initial Portfolio Value	T_0	$1000
Asset 1 Annual Return	$E(\alpha_{t1})$	5%
Asset 2 Annual Return	$E(\alpha_{t2})$	8%
Asset 3 Annual Return	$E(\alpha_{t3})$	12%
Inflation Rate (per annum)	$E(\varphi_{t1})$	3%
Liability Factor Weight, Year 1	$E(\xi_{11})$	0
Liability Factor Weight, Year 2	$E(\xi_{21})$	0.2
Liability Factor Weight, Year 3	$E(\xi_{31})$	0.5

Since our objective is to determine the portfolio value at the end of three years, from equation (1) we get:

$$E(T_3) = ([1+E(\alpha_1)][1+E(\alpha_2)][1+E(\alpha_3)] - E(L_3) - E(L_2)[1+E(\alpha_3)] -$$
$$E(L_1)[1+E(\alpha_2)][1+E(\alpha_3)])T_0$$

since we have only one liability factor return and it has the same value in each time period, $E(\varphi_{11}) = E(\varphi_{21}) = E(\varphi_{31})$, we can reduce the liability terms to;

$$E(L_2) = E(\xi_{21})(1+E(\varphi_{11}))^2$$

$$E(L_3) = E(\xi_{31})(1+E(\varphi_{11}))^3$$

Since $E(\alpha_1) = E(\alpha_2) = E(\alpha_3)$ we can also reduce the asset terms, and substituting;

$$E(\alpha_t) = \omega_1 E(\alpha_{t1}) + \omega_2 E(\alpha_{t2}) + ... + \omega_n E(\alpha_{tn})$$

we obtain,

$$E(T_3) = ([1+\omega_1 E(\alpha_{11}) + \omega_2 E(\alpha_{12}) + \omega_3 E(\alpha_{13})]^3 - E(\xi_{31})[1+E(\varphi_{11})]^3$$
$$- E(\xi_{21})[1+E(\varphi_{11})]^2 [1+\omega_1 E(\alpha_{11}) + \omega_2 E(\alpha_{12}) + \omega_3 E(\alpha_{13})]) * T_0$$

Inserting the terms from Table 2, we have,

$$E(T_3) = ([1+\omega_1(.05) + \omega_2(.08) + \omega_3(.12)]^3 - (.5)[1+(.03)]^3$$
$$-(.2)[1+(.03)]^2 [1+\omega_1(.05) + \omega_2(.08) + \omega_3(.12)]) * 1000$$

If assets 1,2 and 3 are assigned weights of 20%, 50% and 30% respectively, then the expected value of the portfolio at the end of period three is:

$E(T_3)$= $504