

Expected Utility Asset Allocation

William F. Sharpe¹

September, 2006, Revised June 2007

Asset Allocation

Many institutional investors periodically adopt an *asset allocation policy* that specifies target percentages of value for each of several asset classes. Typically a policy is set by a fund's board after evaluating the implications of a set of alternative policies. The staff is then instructed to implement the policy, usually by maintaining the actual allocation to each asset class within a specified range around the policy target level. Such asset allocation (or asset/liability) studies are usually conducted every one to three years or sooner when market conditions change radically.

Most asset allocation studies include at least some analyses that utilize standard mean/variance optimization procedures and incorporate at least some of the aspects of equilibrium asset pricing theory based on mean/variance assumptions (typically, a standard version of the Capital Asset Pricing Model, possibly augmented by assumptions about asset mispricing.)

In a complete asset allocation study a fund's staff (often with the help of consultants) typically:

1. Selects desired asset classes and representative benchmark indices,
2. Chooses a representative historic period and obtains returns for the asset classes,
3. Computes historic asset average returns, standard deviations and correlations
4. Estimates future expected returns, standard deviations and correlations. Historic data are typically used, with possible modifications, for standard deviations and correlations. Expected returns are often based more on current market conditions and/or typical relationships in capital markets
5. Finds several mean/variance efficient asset mixes for alternative levels of risk tolerance,
6. Projects future outcomes for the selected asset mixes, often over many years,
7. Presents to the board relevant summary measures of future outcomes for each of the selected asset mixes, then
8. Asks the board to choose one of the candidate asset mixes to be the asset allocation policy, based on their views concerning the relevant measures of future outcomes.

¹ I am grateful to John Watson of Financial Engines, Inc. and Jesse Phillips of the University of California for careful reviews of earlier drafts and a number of helpful comments.

The focus of this paper is on the key analytic tools employed in steps 4 and 5. In step 5 analysts typically utilize a technique termed *portfolio optimization*. To provide reasonable inputs for such optimization, analysts often rely on informal methods but in some cases utilize a technique termed *reverse portfolio optimization*.

For expository purposes I begin with a discussion of portfolio optimization methods, then turn to reverse optimization procedures. In each case I review the standard analytic approach based on mean/variance assumptions and then describe a more general procedure that assumes investors seek to maximize expected utility. Importantly, mean/variance procedures are special cases of the more general expected utility formulations. For clarity I term the more general approaches *Expected Utility Optimization* and *Expected Utility Reverse Optimization* and the traditional methods *Mean/Variance Optimization* and *Mean/Variance Reverse Optimization*.

The prescription to select a portfolio that maximizes an investor's expected utility is hardly new. Nor are applications in the area of asset allocation. Particularly relevant in this respect is the recent work by [Cremers, Kritzman and Page 2005] and [Adler and Kritzman 2007] in which a "full-scale optimization" numerical search algorithm is used to find an asset allocation that maximizes expected utility under a variety of assumptions about investor preferences.

This paper adds to the existing literature in three ways. First, it presents a new optimization algorithm for efficiently maximizing expected utility in an asset allocation setting. Second it provides a straightforward reverse optimization procedure that adjusts a set of possible future asset returns to incorporate information contained in current asset market values. Finally, it shows that traditional mean/variance procedures for both optimization and reverse optimization are obtained if the new procedures are utilized with the assumption that all investors have quadratic utility.

In large part this paper applies and extends material covered in the author's book *Investors and Markets: Portfolio Choices, Asset Prices and Investment Advice* (Princeton University Press 2006), to which readers interested in more detail are referred.

Mean/Variance Analysis

Much of modern investment theory and practice builds on Markowitz' assumption that in many cases an investor can be concerned solely with the mean and variance of the probability distribution of his or her portfolio return over a specified future period. Given this, only portfolios that provide the maximum mean (expected return) for given variance of return (or standard deviation of return) warrant consideration. A representative set of such mean/variance efficient portfolios of asset classes can then be considered in an asset allocation study, with the one chosen that best meets the board's preferences in terms of the range of relevant future outcomes over one or more future periods.

A focus on only the mean and variance of portfolio return can be justified in one of three ways. First, if all relevant probability distributions have the same form, mean and variance may be sufficient statistics to identify the full distribution of returns for a portfolio. Second, if an investor wishes to maximize the expected utility of portfolio return and considers utility a quadratic function of portfolio return, only mean/variance efficient portfolios need be considered. Third, it may be that over the range of probability distributions to be evaluated, a quadratic approximation to an investor's true utility function may provide asset allocations that provide expected utility adequately close to that associated with a fully optimal allocation, as argued in [Levy and Markowitz 1979].

Asset allocation studies often explicitly assume that all security and portfolio returns are distributed normally over a single period (for example, a year). If this were the case, the focus on mean/variance analysis would be appropriate, no matter what the form of the investor's utility function. But there is increasing agreement that at least some return distributions are not normally distributed, even over relatively short periods and that explicit attention needs to be given to "tail risk" arising from greater probabilities of extreme outcomes than those associated with normal distributions. Furthermore, there is increasing interest in investment vehicles such as hedge funds that may be intentionally designed to have non-normal distributions and substantial downside tail risk. For these reasons, in at least some cases the first justification for mean/variance analysis as a reasonable approximation to reality may be insufficient.

The second justification may also not always suffice. Quadratic utility functions are characterized by a "satiation level" of return beyond which the investor prefers less return to more – an implausible characterization of the preferences of most investors. To be sure, such functions have a great analytic advantage and may serve as reasonable approximations for some investors' true utility functions. Nonetheless, many investors' preferences may be better represented with a different type of utility function. If this is the case, it can be taken into account not only in choosing an optimal portfolio but also when making predictions about tradeoffs available in the capital markets.

While it is entirely possible that in a given setting mean/variance analyses may provide a sufficient approximation to produce an adequate asset allocations, it would seem prudent to at least conduct an alternative analysis utilizing detailed estimates of possible future returns and the best possible representation of an investor's preferences to evaluate the efficacy of the traditional approach. To facilitate this I present more general approaches to optimization and reverse optimization. Specifically, I will assume that forecasts are made by enumerating an explicit set of discrete possible sets of asset returns over a period of choice, with the probability of each outcome estimated explicitly. Given such a set of forecasts, I first show an optimal portfolio would be chosen using a traditional mean/variance analysis then present the more general approach that can take into account an alternative form of an investor's utility function. Subsequently I show how mean/variance methods are used to obtain forecasts consistent with capital market equilibrium and then present a more general approach that can take into account different aspects of the process by which asset prices are determined. For both optimization and

reverse optimization the expected utility approach can provide the same results as the mean/variance approach if the special assumptions of the latter are employed.

Increased generality does not come without cost. The mean/variance approach can be used with continuous distributions (typically, jointly normal). The procedures I present here require discrete distributions, in which each possible scenario is included. At the very least this requires that many alternative outcomes be enumerated. I leave for future research an analysis of practical aspects of this approach, limiting this paper to the presentation of the analytic approaches. Suffice it to say here that the potential advantages of the discrete approach include the ability to allow for more complex distributions than those associated with joint normality and to take into account more aspects of investor preferences.

Mean/Variance Asset Optimization

As shown by [Markowitz 1959] quadratic programming algorithms can be utilized to find the portfolio that provides the maximum expected return for a given level of standard deviation of return. Solving such a quadratic programming problem for each of several different levels of standard deviation of return or variance (standard deviation squared) can provide a set of mean/variance efficient portfolios for use in an asset allocation study.

A standard way to generate a mean/variance efficient portfolio is to maximize a function of expected return and standard deviation of return of the form $d = e - v/t$ where d represents the desirability of the portfolio for the investor, e is the portfolio's expected return, v is its variance of return and t is the investor's risk tolerance. Solving such a problem for different levels of risk tolerance provides a set of mean/variance efficient portfolios for an asset allocation study. By choosing one portfolio from a candidate set of such portfolios the board, in effect, reveals its risk tolerance.

For any given risk tolerance, a mean/variance optimization requires the following inputs:

1. Forecasts of asset return standard deviations
2. Forecasts of correlations among asset returns
3. Expected asset returns
4. Any relevant constraints on asset holdings
5. The investor's risk tolerance

In practice, constraints are often incorporated in such analyses to avoid obtaining "unreasonable" or "infeasible" asset mixes. In many cases inclusion of such constraints reflects inadequate attention to insuring that the asset forecasts used in the analysis are reasonable, taken as a whole.

While general quadratic programming algorithms or other procedures for handling non-linear problems may be used to perform mean/variance optimization, problems in which the only constraints are bounds on the holdings of individual assets can be solved using a simpler gradient method such as that of [Sharpe 1987]. An initial feasible portfolio is analyzed to find the best asset that could be purchased and the best asset that could be sold, where “best” refers to the effect of a small change in holding on the desirability of the portfolio for the investor. As long as the purchase of the former security financed by the sale of the latter will increase the desirability of the portfolio, such a *swap* is desirable. Next, the amount of such a swap to undertake is chosen so as to maximize the increase in portfolio’s desirability, subject to constraints on feasibility. The process is then repeated until the best possible swap cannot increase the portfolio’s desirability. As I will show, a similar approach can be used in a more general setting.

To illustrate both types of optimization I utilize a very simple example. There are three assets (cash, bonds and stocks) and four possible future *states* of the world (alternatively, *scenarios*). Based on history, current conditions and equilibrium considerations, an analyst has produced the forecasts in Table 1². Each entry shows the total return per dollar invested for a specific asset if and only if the associated state of the world occurs.

Table 1
Future Returns

	Cash	Bond	Stock
state1	1.0500	1.0388	0.8348
state2	1.0500	0.9888	1.0848
state3	1.0500	1.0888	1.2348
state4	1.0500	1.1388	1.2848

The states are considered to be equally probable, as shown in Table 2.

Table 2
Probabilities of States

	Probability
state1	0.25
state2	0.25
state3	0.25
state4	0.25

The investor’s risk tolerance (t) equals 0.70. The goal is thus to maximize $e - v/0.70$.

The first step in a mean/variance optimization is to compute the expected returns, standard deviations and correlations of the assets from the future returns (Table 1) and probabilities of states (Table 2). Tables 3 and 4 show the results.

² The numbers in all the tables in this paper are rounded. However, calculations were performed using the original values. The numbers in Table 1 were produced using reverse optimization. For details, see the subsequent discussion and the description of Table 16.

Table 3
Expected returns and standard deviations

	E	SD
Cash	1.0500	0.0000
Bond	1.0638	0.0559
Stock	1.1098	0.1750

Table 4
Correlations

	Cash	Bond	Stock
Cash	1.0000	0.0000	0.0000
Bond	0.0000	1.0000	0.6389
Stock	0.0000	0.6389	1.0000

In this case there are no constraints on asset holdings other than that the sum of the proportions invested in the assets equals one. The resulting mean/variance optimal portfolio is shown in Table 5.

Table 5
Optimal portfolio

Cash	0.0705
Bond	0.3098
Stock	0.6196

Had there been lower bounds on the asset holdings of zero and upper bounds of one, the same portfolio would have been obtained since the constraints would not have been binding.

Expected Utility Optimization

I now introduce the more general approach to optimization and illustrate its use with our simple example. I will show that the procedure will produce the same results as mean/variance optimization if the investor is assumed to have a particular type of preferences, but that a different type of preferences will give a different optimal portfolio.

Expected Utility

The key assumption is that the goal of an investor is to maximize the *expected utility* of the return from his or her portfolio. Associated with the portfolio return in each state of the world is a *utility* which measures the “happiness” associated with the total return in that state. The expected utility of the return in a state equals its utility times the probability that the state will occur. The expected utility of the portfolio is then the sum of the expected utilities of its returns in the states.

The utility of a total portfolio return R_{ps} in state s will be denoted $u(R_{ps})$. The investor's goal is to maximize expected utility eu which equals³:

$$(1) \quad eu = \sum_s \pi_s u(R_{ps})$$

Where π_s is the probability that state s will occur⁴.

Note that this assumes that the utility function is the same for all states and that expected utility is separable and additive across states. These assumptions rule out some aspects of preferences but are considerably more general than the mean/variance approach.

The first derivative of utility, *marginal utility* of total portfolio return in state s will be denoted $m(R_{ps})$. I assume that marginal utility decreases with R_{ps} . This is equivalent to assuming that the investor is risk-averse.

I assume that the only constraints are upper and lower bounds on individual asset holdings of the form:

$$(2) \quad lb_i \leq x_i \leq ub_i$$

Where x_i is the proportion of the fund invested in asset i and the x_i values sum to 1.

Maximizing Expected Utility

I now present an algorithm for solving the nonlinear programming problem described in the previous section.

First, note that the marginal expected utility (*meu*) of the portfolio return in state s is:

$$(3) \quad meu(R_{ps}) = \pi_s m(R_{ps})$$

³ In this and some subsequent formulas I use two or more letters for some variable names. This is a convention familiar to those who write computer programs but may cause consternation for some mathematicians. Hopefully the meanings of the formulas will be plain in context.

⁴ The use of the symbol for pi for probability may cause some confusion but the mnemonic value is considerable and the presence of a subscript should serve as a reminder that this is not 3.14159....

Now consider the effect on portfolio expected utility of a small change in the amount invested in asset i . Since \$1 invested in the asset provides R_{i1} in state 1, R_{i2} in state 2, etc.) the marginal expected utility (per dollar) of asset i will be:

$$(4) \quad meu_i = \sum_s R_{is} meu(R_{ps})$$

This will hold for any asset and for any portfolio, since a portfolio's return is simply a weighted sum of asset returns, with the weights summing to 1.

Note that for each asset the cost of obtaining its set of returns across states is \$1. Now consider a portfolio for which there are two assets, i and j , where $meu_i > meu_j$ and it is feasible to purchase additional units of asset i and to sell some units of asset j . Clearly the portfolio can be *improved* – that is, its expected utility can be increased. The reason is straightforward. Since it costs \$1 to obtain each asset's set of returns, if these three conditions are met, expected utility can be increased by selling some units of asset j and using the proceeds to buy units of asset i .

It is a simple matter to determine whether or not such a change is possible for a given portfolio. The marginal expected utility for each asset can be computed and each asset classified as a potential buy (if $x_i < ub_j$) and/or a potential sell (if $x_i > lb_j$). The *best buy* is the asset among the potential buys with the largest marginal expected utility. The *best sell* is the one among the potential sells with the smallest marginal expected utility. If the marginal expected utility of the best buy exceeds that of the best sell, then the best swap involves selling units of the best sell and purchasing units of the best buy. If this is not the case, or if there are no potential buys or no potential sells, then the portfolio cannot be improved.

Once a desirable swap has been identified, the optimal magnitude for the amount to be swapped can be determined. A simple procedure determines the largest feasible magnitude, given by the upper bound of the asset to be bought and the lower bound of the asset to be purchased. The marginal expected utilities of the assets that would obtain were this swap undertaken are then determined. If the spread between the marginal utility of the asset being bought and that being sold would still be positive, the maximum swap should be made. Otherwise, an intermediate amount (for example, half way between the minimum of zero and the maximum) should be considered and the marginal expected utilities that would obtain were that swap undertaken calculated. If the spread would still be positive, the range of swaps should be restricted to that between the intermediate amount and the maximum amount. Otherwise the range should be restricted to that between the minimum and the intermediate amount. This procedure is then continued until the marginal expected utility spread is lower than a desired threshold or a similar condition is met for the difference between the current maximum and minimum swap amounts.

The overall algorithm for finding the maximum expected utility portfolio is thus:

1. Find a feasible portfolio
2. Determine the best possible two-asset swap
3. If no such swap is feasible, terminate
4. Otherwise, compute the best magnitude for the swap and revise the portfolio
5. Repeat starting at step 2 until the condition in step 3 is met.

Up to the precision utilized for the termination conditions, this procedure will find a portfolio that maximizes expected utility under the specified conditions, that is only upper and lower bounds on asset holdings and marginal utility that decreases with portfolio return.

Quadratic Utility Functions

I first illustrate the use of our optimization procedure with a case in which the investor's utility is a quadratic function of portfolio return. Moreover, I choose an investor with an attitude towards risk that is equal to that of the investor considered earlier.

An investor with quadratic utility can be described as maximizing the expected value of a quadratic function of portfolio return of the form:

$$(5) \quad u(R) = R - \frac{k}{2}R^2$$

The associated marginal utility function is:

$$(6) \quad m(R) = 1 - kR$$

The satiation level is that for which marginal utility equals $1/k$; since for higher levels of return the marginal utility will be negative and total utility will be lower than that associated with a return of $1/k$. The higher is k , the more risk an investor with quadratic utility will take, other things equal.

For an investor with quadratic utility, expected utility will be a linear function of the expected portfolio return and the expected value of the return squared. But the latter is equal to the portfolio variance plus the expected return squared. Thus expected utility will be a function of portfolio expected return and variance of return (although not of the linear form used in the mean/variance optimization procedure described earlier). For any given expected return an investor with quadratic utility will wish to minimize variance and hence will choose a mean/variance efficient portfolio.

Now consider a quadratic utility investor with $k=0.6938$ (to four decimal places) and a corresponding satiation return of 1.4413. The expected utility maximization algorithm gives the optimal portfolio shown in Table 6.

	Portfolio
Cash	0.0705
Bond	0.3098
Stock	0.6196

This is precisely the same portfolio shown in Table 5, which was obtained using the mean/variance optimization quadratic programming optimization procedure. This is because the quadratic utility parameter k was set to a value determined by the riskless rate of return (r) and the expected return (e) and variance (v) of the optimal portfolio found earlier, as follows⁵:

$$(7) \quad k = \frac{e-r}{v+e^2-re}$$

In this case, the same portfolio is optimal for an investor maximizing $e-v/t$ where $t=0.70$ and one maximizing expected utility where utility equals $R-(k/2)R^2$ and $k=0.6938$. A higher value of k would give a portfolio with a greater mean and variance while a lower value of k would give a more conservative portfolio. In any event the algorithm will produce a mean/variance efficient portfolio whenever the utility function is quadratic. Thus, as claimed earlier, mean/variance optimization can be considered a special case of the more general expected utility optimization approach.

HARA Utility Functions

While the optimization algorithm can be used for an investor with a quadratic utility function, its main advantage is the ability to find an optimal portfolio for an investor with another type of utility function. The only requirements are that the expected utility of the portfolio is a separable additive function of the expected utilities in the alternative states, that marginal utility is a decreasing function of portfolio return and that the only constraints are upper and lower bounds on asset holdings.

To illustrate, consider a function of considerable generality, known as a HARA function (since an investor with such a function exhibits *hyperbolic absolute risk aversion*). More specifically, the investor's utility is related to portfolio return as follows:

$$(8) \quad u(R) = \frac{(R-b)^{1-c}}{1-c}$$

⁵ This relationship can be derived from the characteristics of the two optimization problems.

The value of b can be considered the investor's *minimum required level return* since the marginal utility of increasing return above this level is infinite. Thus only values of R greater than b should be considered. Above this level the investor's marginal utility will be related to portfolio return by:

$$(9) \quad m(R) = (R - b)^{-c}$$

In most cases $c > 1$ although the marginal utility equation (which is the only one needed for our algorithm) may be used when $c = 1$ (and utility is equal to the logarithm of return) or when c is less than one and positive.

An important special case arises when $b = 0$. In such a case the investor is said to have *constant relative risk aversion (CRRA)*. For such an investor a given percentage increase in return will cause a given percentage decrease in marginal utility, regardless of the initial level of return. More generally, a given percentage increase in return over and above the minimum required level ($R - b$) will cause a given percentage decrease in marginal utility, regardless of the initial level of return. With respect to total return (R), an investor with $b > 0$ has decreasing relative risk aversion, while one with $b < 0$ has increasing relative risk aversion.

The first column in Table 7 shows the results obtained using the data in Tables 1 and 2 for an investor with constant relative risk aversion ($b = 0$) and a risk aversion coefficient (c) of 4. Upper and lower bounds on asset holdings were set to allow short positions in any asset but none were binding. As can be seen, the optimal portfolio includes long holdings in all three assets, with a larger holding in bonds than in stocks. The table also shows the returns on this portfolio in the four states of the world as well as key characteristics of the portfolio.

Portfolio:	EU Opt	MV Opt
Cash	0.0885	0.2956
Bond	0.5101	0.2347
Stock	0.4014	0.4698
Returns:		
state1	0.9579	0.9463
state2	1.0327	1.0520
state3	1.1440	1.1459
state4	1.1896	1.1811
Characteristics:		
CER	1.0654	1.0651
ExpRet	1.0810	1.0813
SDRet	0.0911	0.0911
SR	0.3406	0.3436

The first characteristic is the *certainty-equivalent return* (CER) of this portfolio for the investor in question. The CER is a return which, if obtained in every state, would provide the same expected utility as the portfolio itself. As can be seen, this portfolio is as desirable for this investor as a return of 1.0654, which is considerably greater than the riskless rate of interest (1.050). Importantly, no other feasible portfolio can provide a greater CER for this investor.

The remaining rows in Table 7 show conventional statistics for the max-EU portfolio: its expected return, standard deviation of return and Sharpe Ratio (the ratio of expected return minus the riskless rate to standard deviation).

The second column in Table 7 shows the characteristics of a portfolio obtained with mean/variance optimization chosen to have the same standard deviation of return. As can be seen, its composition is considerably different, with a substantially smaller position in bonds than in stocks.

Table 7 also shows that the two approaches to optimization provide quite different returns in the states. In particular, the mean/variance portfolio has a lower return in the worst state, a lower return in the best state and higher returns in the intermediate states. By design it has the same standard deviation of return as the Max-EU portfolio and, given the manner of its derivation, a higher expected return and Sharpe Ratio. But it is nonetheless inferior in terms of the expected utility for the investor in question, with a certainty equivalent return of 1.0651, three basis points smaller than the maximum possible CER of 1.0654.

Clearly, a mean/variance optimal portfolio may not be optimal for an investor with non-quadratic utility. Such an investor may willingly accept a portfolio with a lower Sharpe Ratio than is available at the same risk level as conventionally measured (by standard deviation of return). Similarly, an investor with quadratic utility may choose a portfolio with a lower certainty-equivalent return as measured by an investor with some other type of utility function. Given diverse investor preferences there should be diverse portfolio holdings. The expected utility optimization procedure makes it possible to find preferred asset holdings for investors whose utility functions differ in form as well as in parameters.

Reverse Optimization

Whether a mean/variance or expected utility approach is used for asset allocation optimization, the quality of the result depends crucially on the meaningfulness of the inputs. Such inputs should at the very least take into account historic returns, current asset market values and the best possible assumptions about relationships among returns in capital markets.

Formal procedures for doing this have been developed for use in mean/variance analyses. Some analysts use the procedures explicitly; others adopt judgmental approaches motivated by such considerations. This section begins by describing the standard reverse optimization procedure applied in a mean/variance context. Next I provide a more general approach using expected utility and show that with the assumption of quadratic utility the same results can be obtained as with the standard approach. Finally I show that the more general approach can be used to reflect more general assumptions about equilibrium in capital markets.

Key to all these approaches is the goal that a set of return forecasts conform with or be based on a view of the nature of equilibrium in capital markets and with the current market values of the assets in question. In many cases, asset allocation studies are intended to find asset mixes that would be appropriate assuming that the markets for major asset classes are *informationally efficient* in the sense that such prices fully reflect available information about the uncertain future prospects for those asset classes. Such forecasts are sometimes termed *efficient market forecasts* or *passive forecasts*. Investors who aspire to choose an asset allocation based on efficient market assumptions assume that asset classes are priced reasonably, although they may assume that within asset classes some securities are mispriced.

In some cases analysts prefer to utilize forecasts for an asset allocation study that incorporate an assumption that some or all asset classes are mispriced. However, such an approach requires an underlying set of forecasts consistent with correct pricing of all asset classes. We concentrate here on the derivation of such efficient market forecasts, leaving for later research the subject of possible adjustments to reflect perceived mispricing.

Mean/Variance Reverse Optimization

The standard mean/variance approach to reverse optimization assumes that the conditions of the Capital Asset Pricing Model hold, that is that there is a particular relationship between each asset's expected return and a measure based on asset risks, correlations and the current market values of the asset classes.

In many cases, practitioners utilize historic asset standard deviations and correlations either directly or with some modification as predictions of the corresponding future risks and correlations. Asset expected returns are then determined from current asset market values and the conditions of the CAPM.

The CAPM implies that in equilibrium, each asset's expected excess return over the riskless rate of interest is proportional to the asset's *beta* value, which is in turn based on its covariance with the *market portfolio* of all asset classes. Given a set of asset standard deviations and correlations and the current market values of the asset classes, it is a simple matter to compute each asset's beta value. Then, given an assumed expected return for the market portfolio, each asset's expected return can be computed. The resultant set of return forecasts can then be considered efficient market forecasts.

This process is equivalent to solving a mean/variance optimization process in reverse, hence the name. As we have seen, a mean/variance optimization problem is of the form:

Expected Returns + Risks + Correlations + Risk Tolerance → Optimal Portfolio

In this case it is assumed that the market portfolio is optimal for a “representative investor” with a particular risk tolerance. Risks and correlations are known, so the conditions for portfolio optimality are used to find the set of expected returns that would give the market portfolio as a solution to a mean/variance optimization problem. Thus:

Risks + Correlations + Representative Risk Tolerance + Market Portfolio → Expected Returns

The use of reverse optimization for obtaining inputs for asset allocation studies was advocated in [Sharpe 1985]. Subsequently, [Black and Litterman 1992] extended this approach to show how views about asset mispricing can be utilized to modify the initial reverse optimization results. In each case *current* market values of the asset classes are used to insure that the forecasts taken as a whole are plausible.

Not every asset allocation study includes reverse optimization based on current asset market values. This accounts in large part for the relatively frequent use of *ad hoc* constraints or other methods to insure plausibility for the results obtained in the optimization phase. This is unfortunate, for current market values represent highly useful information about investors' collective assessments of the present values of the uncertain future prospects of asset classes. It is foolish to fail to take this important information into account when making crucial policy decisions.

To illustrate mean/variance reverse optimization I use the analysis that produced the forecasts in Tables 1 and 2.

Table 8 shows the historic returns for the three assets in four prior years. The return on the first asset in each year can be thought of as the yield on a one-year Treasury bill purchased at the beginning of the year. The second and third can be considered the returns on index funds of bonds and stocks, respectively. The table entries show the actual returns for each asset in each year in the historic period although they are labeled state1, state2, etc. to reflect the fact that the entries will be transformed to reflect outcomes in alternative possible future states of the world. To keep the analysis simple, each state is assumed to be equally probable.

Table 8
Historic Returns:

	Cash	Bond	Stock
state1	1.0600	1.0200	0.7200
state2	1.0700	0.9800	0.9800
state3	1.0600	1.0700	1.1200
state4	1.0500	1.1100	1.1600

Table 9 shows the relative values for the assets at the end of the most recent year.

Table 9
Current Market Portfolio:

	Cash	Bond	Stock
Rel. Value	0.10	0.30	0.60

The information in Table 8 is based entirely on historic information. The information in Table 9 uses current data to reflect market participants' current forecasts of the assets' future prospects.

The remaining information needed for the reverse optimization procedure also concerns the future. It is shown in Table 10. The first entry is the current riskless rate of interest and the second an estimate of the difference between the expected return of the market portfolio and the current riskless rate (often termed the *market risk premium*).

Table 10
Current Forecasts:

	Value
Riskless Rate	1.050
Mkt Risk Prem.	0.040

The goal of reverse optimization is to combine the information in Tables 8, 9 and 10 to obtain relevant forecasts of possible future returns.

The first step is to convert the returns in Table 8 to *excess returns* by subtracting from each one the return on the riskless asset in the year in question. The excess returns are then added to the current riskless rate of interest to obtain forecasts more consistent with current observed market information. This gives the returns in Table 11.

Table 11:
Excess Returns plus Current Riskless Rate:

	Cash	Bond	Stock
state1	1.0500	1.0100	0.7100
state2	1.0500	0.9600	0.9600
state3	1.0500	1.0600	1.1100
state4	1.0500	1.1100	1.1600

The next step is to compute the projected covariances of the assets. Since the covariance of assets i and j equals their correlation times the product of their standard deviations this incorporates both the historic risks and correlations of the assets' excess returns. The covariance matrix is shown in Table 12.

Table 12
Covariance Matrix:

	Cash	Bond	Stock
Cash	0.000000	0.000000	0.000000
Bond	0.000000	0.003125	0.006250
Stock	0.000000	0.006250	0.030625

As shown earlier, the standard deviations are 0, 0.0559 and 0.1750 for cash, bonds and stocks, respectively and the correlation between bonds and stocks is 0.6389.

To compute the beta of every security relative to the current market portfolio requires combining the information in Table 12 with that in Table 9. The beta value for an asset is its covariance with the market portfolio divided by the variance of the market portfolio. The covariance of an asset with the market portfolio is a weighted average of its covariances with the assets, using the market portfolio proportions as weights. The variance of the market portfolio is the weighted average of the asset covariances with the market using the market proportions as weights. These formulas give the beta values shown in Table 13.

Table 13
Beta Values:

	Beta
Cash	0.0000
Bond	0.3458
Stock	1.4938

Given the beta values, the market risk premium and the riskless rate of interest, it is straightforward to compute a set of security expected returns consistent with the CAPM. The expected return on a security is simply the riskless rate of interest plus its beta value times the market risk premium. Table 14 shows the results of these calculations.

Table 14
Expected returns:

	ExpRet
Cash	1.0500
Bond	1.0638
Stock	1.1098

This completes the computation of efficient market forecasts for a standard mean/variance optimization. The information required for the optimization phase is contained in Tables 12 and 14. Importantly, these parameters take into account historic information about asset return variability and covariability, the current market values of the asset classes, the current riskless rate of interest, the estimated future market risk premium and an assumed set of relationships that would obtain in a market conforming with the CAPM.

The covariance matrix in Table 12 and the expected returns in Table 14 were the inputs used for the mean/variance optimization analysis in the earlier section of this paper. The optimal portfolio in Table 5 can thus be considered the appropriate strategy to be followed by a mean/variance investor with a risk tolerance (t) of 0.70 operating in a market in which the market portfolio is mean/variance efficient and the risk premium is 4%.

The Representative Investor

When optimization inputs are created using mean/variance reverse optimization the solution for one specific level of risk tolerance will be the market portfolio. The risk tolerance t in the objective function $e - v/t$ for which the market portfolio is optimal will equal twice the market portfolio's variance divided by the market risk premium. In this case the resulting "market risk tolerance" equals 0.6778. In effect, the market prices of the assets are those that would prevail if there were a single investor with quadratic utility who wished to maximize $[e - v/0.6778]$. Such an investor is often termed a *representative investor*. In a world of unconstrained investors with quadratic utility this is a reasonable interpretation, with the market risk tolerance equal to a wealth-weighted average of the risk tolerances of the actual investors. In a more realistic setting, however, the relationship is likely to be more complex. Nonetheless, it is useful to characterize asset returns as consistent with a capital market with a single representative investor of a particular type.

While we will use the term "representative investor" it is important to emphasize that the key ingredient in all that follows is the use of a marginal utility function consistent with market equilibrium. This will reflect the marginal utility functions of actual investor but may not be a simple weighted average of their individual marginal utility functions. A more appropriate term for the market function in question is the *asset pricing kernel*, although for expository purposes I will continue to refer to it as the marginal utility function of the representative investor. For an in-depth discussion of this issue see [Sharpe 2007].

Obtaining Return Forecasts Using Reverse Optimization

As before I assume that optimization will be based on a discrete set of possible future asset returns in alternative future states of the world and the probabilities associated with those states. The goal is to develop a reverse optimization procedure that can produce such a set of forecasts. I show first how this can be accomplished in a standard mean/variance setting, then provide a more general approach.

Obtaining Return Forecasts using Mean/Variance Reverse Optimization

In the previous example, the standard mean/variance reverse optimization procedure used the returns in Table 11 to generate the covariance matrix in Table 12. Then the covariance matrix and the current market portfolio were used to determine asset beta values. Finally, the beta values were combined with the riskless return and the market risk premium to determine the expected returns shown in Table 14. The present goal is to provide a procedure that will produce a return table and set of probabilities consistent with the covariance matrix in Table 12 and the expected returns in Table 14. The approach is quite simple. Adding a (possibly zero or negative) constant to the return on an asset in every state will not affect its covariance with any other asset. Thus for each asset class a constant can be selected that will give the desired expected return without altering the asset risks, correlations and hence covariances.

The goal is to select a set of differences (d_i) such that the new return (R_{is}) for each asset in each state is equal to the old return (R_{is}^o) plus the selected difference:

$$(10) \quad R_{is} = R_{is}^o + d_i$$

Taking expected values over states gives:

$$(11) \quad E(R_i) = E(R_i^o) + d_i$$

In this case the old returns are those shown in Table 11. The first column in table 15 shows the expected returns for the assets based on those returns, the second column the desired expected returns (from Table 14) and the third column the constant (d_i) to be added to each asset's returns to give the desired expected return.

Table 15
Constants for New Returns

	Old ER	New ER	Difference
Cash	1.0500	1.0500	0.0000
Bond	1.0350	1.0638	0.0288
Stock	0.9850	1.1098	0.1248

Table 16 shows the set of returns obtained by adding these constants⁶ to the adjusted returns, that is, the excess returns plus the risk-free rate. As intended, this will have the same covariance matrix as the returns in Table 11 and the desired expected returns from Table 14.

Table 16
New Security Returns

	Cash	Bond	Stock
state1	1.0500	1.0388	0.8348
state2	1.0500	0.9888	1.0848
state3	1.0500	1.0888	1.2348
state4	1.0500	1.1388	1.2848

Note that Table 16 is the same as Table 1. In fact, the latter is the origin of the former, used for the optimization examples. The forecasted returns for the assets in alternative states were produced in the manner shown here using historic returns, the current riskless return, an assumption about the expected return on the market, the current values of the assets and an assumption that equilibrium in the capital markets obtains and is consistent with a representative investor who cares only about portfolio mean and variance. They are thus efficient market forecasts with no alterations to reflect possible mispricing.

Expected Utility Reverse Optimization

I now turn to a more general approach to reverse optimization. To preserve the information about return variation and covariation contained in historic returns, I take as a constraint the requirement that each new return for an asset in a state equal the old return in that state plus a constant. However, the constants must be consistent with the requirement that the market portfolio provide the maximum expected utility for the representative investor.

Note first that the market portfolio's return in a state using the desired new returns will equal its return in that state using the old returns plus the market value-weighted sum of the asset constants:

$$(12) \quad R_{ms} = \sum_i x_{im} R_{is} = \sum_i x_{im} R_{is}^o + \sum_i x_{im} d_i$$

⁶ As indicated in the footnote to Table 1, values shown in this paper are rounded. For those wishing to replicate our results, the constants added to bond and stock to produce table 16 were 0.02883125864454 and 0.1247510373444, respectively.

For each state, the new return on the market will equal the old return plus a constant (d_m) which will, in turn equal a value-weighted sum of the asset constants:

$$(13) \quad d_m = \sum_i x_{im} d_i$$

But this value must be set to produce the desired expected return on the market portfolio - the sum of the riskfree rate and the market risk premium, that is:

$$(14) \quad E(R_m) = E(R_m^o) + d_m$$

Table 17 shows the old and new values and the constant to be added to each old market return.

Table 17
Market Expected Returns:

	Old	New	d_m
ExpRet	1.0065	1.0900	0.0835

Adding 0.0835 to every market return gives the new market returns shown in Table 18.

Table 18
Market Returns:

	Return
state1	0.9175
state2	1.0525
state3	1.1725
state4	1.2175

I now turn to the conditions required for a portfolio to maximize expected utility. Recall first that the marginal expected utility of asset i will be related to its returns in the states and the marginal expected utilities for the states. For an investor who holds the market portfolio, the marginal expected utility of that portfolio will be:

$$(15) \quad meu_m = \sum_s R_{ms} \pi_s m(R_{ms})$$

Moreover for such an investor, the marginal expected utility of the riskfree asset, which returns R_f in every state will be:

$$(16) \quad meu_f = R_f \sum_s \pi_s m(R_{ms})$$

But each of these assets sells for the same price (\$1). Hence, for the market portfolio to be optimal, their marginal expected utilities must be the same:

$$(17) \quad \sum_s R_{ms} \pi_s m(R_{ms}) = R_f \sum_s \pi_s m(R_{ms})$$

All the values in this equation are given; only the form and parameters of the representative investor's marginal utility function are free to be determined. I illustrate the manner in which this can be done in two cases.

Expected Utility Reverse Optimization with Quadratic Utility

In this section I assume that the representative investor's utility function is quadratic; thus the marginal utility function is linear. That is:

$$(18) \quad m(R_{ms}) = 1 - kR_{ms}$$

Substituting equation (18) in equation (17) and simplifying gives:

$$(19) \quad E(R_m) - kE(R_m^2) = R_f - kR_fE(R_m)$$

For this equation to hold, the value of k must satisfy:

$$(20) \quad k = \frac{E(R_m) - R_f}{E(R_m^2) - R_fE(R_m)}$$

Using the market returns from Table 18 gives a value of k equal to 0.6998. The market portfolio will be optimal for a quadratic utility investor with this value of k and hence a satiation return level (1/k) of 1.4289.

It remains to determine the constants to be added to the individual asset returns. This is easily accomplished. Since every asset costs \$1 it must be the case that the marginal expected utility for each asset must be the same. For example, the marginal utility of asset i must equal that of the riskless asset, that is:

$$(21) \quad \sum_s R_{is} \pi_s m(R_{ms}) = R_f \sum_s \pi_s m(R_{ms})$$

But the return on asset i in state s is its old return plus d_i :

$$(22) \quad R_{is} = R_{is}^o + d_i$$

Thus d_i must be set so that:

$$(23) \quad \sum_s (R_{is}^o + d_i) \pi_s m(R_{ms}) = R_f \sum_s \pi_s m(R_{ms})$$

The required value of d_i is thus:

$$(24) \quad d_i = \frac{R_f \sum_s \pi_s m(R_{ms}) - \sum_s R_{is}^o \pi_s m(R_{ms})}{\sum_s \pi_s m(R_{ms})}$$

Table 19 shows the resulting d_i values.

Table 19	
Return Differences:	
	Difference
Cash	0.0000
Bond	0.0288
Stock	0.1248

The market value-weighted average of the d_i values in Table 19 equals 0.0835, the value computed earlier to be added to each market return to obtain the desired expected return on the market portfolio. The two procedures will always give the same value, whether the utility function is quadratic or not, as can be seen by multiplying either of the previous two equations by x_i and summing over the assets.

Table 20 shows the resulting returns table. Not surprisingly, it is the same as the ones shown in Tables 16 and 1. Thus this general approach gives the same results as the standard mean/variance optimization when the representative investor's utility function is quadratic.

Table 20
New Returns:

	Cash	Bond	Stock
state1	1.0500	1.0388	0.8348
state2	1.0500	0.9888	1.0848
state3	1.0500	1.0888	1.2348
state4	1.0500	1.1388	1.2848

Expected Utility Reverse Optimization with HARA Utility

I now consider the second case. Here the representative investor is characterized by a HARA utility function. Thus:

$$(25) \quad m(R_{ms}) = (R_{ms} - b)^{-c}$$

As before the marginal expected utility of the market portfolio must equal that of the riskless asset:

$$(26) \quad \sum_s R_{ms} \pi_s m(R_{ms}) = R_f \sum_s \pi_s m(R_{ms})$$

Thus the goal is to find values of b and c that satisfy:

$$(27) \quad \sum_s R_{ms} \pi_s [(R_{ms} - b)^{-c}] = R_f \sum_s \pi_s [(R_{ms} - b)^{-c}]$$

There are multiple pairs of such values. To simplify, one can pre-specify the value of b , then solve for c . Since the equation is non-linear, a search over alternative values of c is required, but this can be done quickly using any standard procedure. If desired, the procedure can be repeated for different values of b and the most realistic combination chosen.

For this example I assume that $b=0$ so the representative investor will have constant relative risk aversion (an assumption frequently made when calibrating asset pricing models). Using the inputs in Tables 8, 9 and 10 (as before) this gives a value of 2.9868 for c .

Once the values of b and c have been specified it is straightforward to compute the marginal expected utility for each state. Then the d_i values can be computed using the procedure employed earlier. Table 21 shows the results. As can be seen, they differ from those in Table 19 which were based on a quadratic utility function, but not by large amounts.

Table 21
Security differences:

	Difference
Cash	0.0000
Bond	0.0269
Stock	0.1257

Table 22 shows the forecasted returns, obtained by adding the d_i values to the returns for the assets. Of course these differ from the returns in Table 20 which were based on a quadratic utility function, but not by large amounts.

Table 22
New Security Returns:

	Cash	Bond	Stock
state1	1.0500	1.0369	0.8357
state2	1.0500	0.9869	1.0857
state3	1.0500	1.0869	1.2357
state4	1.0500	1.1369	1.2857

State Prices

In the discussions of both optimization and reverse optimization I have focused on the marginal expected utilities of states and assets. A closely related concept is the rate at which an investor would be willing to substitute one asset or state for another. This will equal the ratio of the associated marginal expected utilities. Of particular interest is mrs_s , the marginal rate of substitution of the riskless asset that pays \$1 in every state for a *state claim* paying \$1 in state s . Letting R_{ps} be the return on the investor's portfolio in state s :

$$(28) \quad mrs_s = \frac{\pi_s m(R_{ps})}{\sum_s \pi_s m(R_{ps})}$$

Since the riskless asset costs \$1 and returns R_f , the marginal rate of substitution of present value for the riskless asset is $1/R_f$. Thus the marginal rate of substitution of present value for a state claim paying \$1 in state s is equal to the product of the two marginal rates of substitution. This can be considered p_s , the investor's reservation price or, more simply, the *state price*, for the state claim:

$$(29) \quad p_s = \frac{\pi_s m(R_{ps})}{R_f \sum_s \pi_s m(R_{ps})}$$

Table 23 shows the state prices for the example. Not surprisingly, they sum to $1/1.05$.

Table 23
State Prices:

	Price
state1	0.3700
state2	0.2456
state3	0.1779
state4	0.1589

In a reverse optimization the reservation state prices can be used as estimates of the values of assets not included explicitly in the analysis. Any asset that is traded should sell for a price close to the value computed from the state prices. This makes it possible to evaluate the extent to which the derived representative investor's utility function provides a reasonable characterization of equilibrium in the capital market and the returns on the assets in it.

To illustrate, assume that there is a traded put option on the stock index fund with an exercise price of \$1. It will be in the money only if state 1 occurs. Assuming that adjustments are made for dividends, it will pay 0.0825 in state 1 and zero in every other state. Its price should thus equal $0.3700 \cdot 0.0825$, or 0.0305. If the actual price differs significantly from this, it may be desirable to perform another reverse optimization with a different type of utility function or a different fixed parameter for the same type of utility function. More generally, the values of traded derivative securities on underlying assets that are included in the reverse optimization can be used to improve the entire procedure. In this example, different values of b could be selected, with the results for each evaluated on the basis of the conformance of the actual prices of asset derivatives to the values computed using the associated state prices.

Applications

This paper has been limited to the presentation of expected utility approaches to asset optimization and reverse optimization. We leave for other research the task of evaluating the practicality of such methods in actual applications. A necessary condition for success in applications is the creation of plausible discrete forecasts of alternative sets of asset returns. If reasonable estimates of such returns can be obtained, the procedures described here could offer a number of advantages. Complex distributions of possible future outcomes could be analyzed. Differences between an investor's views about the future and those reflected in asset prices could be incorporated by adding "alpha values" to asset returns and/or changing the probabilities associated with some or all of the states. A wide range of investor preferences could be considered. The optimization procedure could be extended to cover cases in which assets are to be allocated taking liabilities into account. And the procedures could also be applied in the area of risk management and budgeting.

Perhaps most important, expected utility asset allocation analysis could potentially replace many *ad hoc* procedures currently used by investment staffs and boards, providing a single and consistent analytic approach. The practical challenges may be formidable but the possible gains could be substantial.

References

Adler, Tim and Mark Kritzman, "Mean-variance versus full-scale optimization," *Journal of Asset Management*, 2007, Vol. 7, Number 5.

Black, Fischer, and Robert Litterman. 1992. "Global Portfolio Optimization." *Financial Analysts Journal*, September/October, 28-43.

Cremers, Jan-Hein, Mark Kritzman and Sebastien Page, "Optimal Hedge Fund Allocations: Do Higher Moments Matter?" *The Journal of Portfolio Management*, Spring 2005.

Levy, Haim and Harry M. Markowitz, 1979. "Approximating Expected Utility by a Function of Mean and Variance," *American Economic Review*, June, Vol. 69, No. 3.

Markowitz, H.M. 1959, *Portfolio Selection: Efficient Diversification of Investments*, Wiley, Yale University Press.

Sharpe, William F. 1985. *AAT: Asset Allocation Tools*, Redwood City, Calif: The Scientific Press

Sharpe, William F. 1987. "An Algorithm for Portfolio Improvement," in K.D. Lawrence, J.B. Guerard and Gary D. Reeves, Editors, *Advances in Mathematical Programming and Financial Planning*, JAI Press, Inc., pp. 155-170.

Sharpe, William F. 2007. *Investors and Markets: Portfolio Choices, Asset Prices and Investment Advice*, Princeton: Princeton University Press.