The Assumptions behind Mean-Variance

Increasing expected return while minimizing the variance may sound like a sensible policy, but we need to be more precise about the assumptions that are involved. We start with some basic assumptions about investor behavior.

Choice Axioms

Most economic theories of the choices that individuals make between uncertain outcomes derive from four axioms that were originally proposed by von Neumann and Morgenstern.

Think of the problem of George who is faced with the choice of several portfolios offering different possible outcomes. What can we predict about the choices that George may make? If he acts rationally, his choices will obey the following axioms:

Axiom 1 Completeness (sometimes called Comparability) For any two outcomes, x and y, George can always state that he prefers x to y, or y to x, or that he is indifferent between the two.

Axiom 2 Transitivity (sometimes called Completeness) If George prefers outcome x to y and y to z, then he will also prefer x to z.

Axiom 3 Strong Independence Suppose one portfolio offers a payoff of x with a probability of α and z with a probability of (1 – α). Suppose also that George is offered an alternative portfolio which offers a payoff of y with a probability of α and a mutually exclusive payoff of z with a probability of (1 – α). If George is indifferent between x and y, then he must be indifferent between the two portfolios.

Axiom 4 Continuity (sometimes called Measurability) If George prefers outcome x to y and y to z, then there is some probability α that he will be indifferent between y and a portfolio that offers x with probability α and z with probability (1 – α).

Expected Utility

Suppose now that George has $1000 to invest and needs to choose between two packages of investments. Portfolio A would provide him with $1,200 with a probability of .6 and -$900 with a probability of .4. Its expected payoff is (.6 x 1,200) + (.4 x 900) = $1,080, an expected return of 8%. Portfolio B offers a payoff of $1,140 with a probability of .5 and $1,000 also with a probability of .5. Its expected payoff is (.5 x 1,400) + (.5 x 1,000) = $1,070, an expected return of 7%. Would George be acting irrationally if he chose Portfolio B with the lower expected return? Not necessarily, for B has a lower spread of possible outcomes. If George is risk-averse, he may well prefer B. It provides a lower expected return but it may provide a higher *expected utility.*

If George's behavior is consistent with the four axioms of rational behavior, then there exists some *utility* that he derives from each possible outcome, and his decisions can be characterized as if he were maximizing expected utility.

Suppose that we arbitrarily assign 100 units of utility to $1000 of wealth but only 80 units to $900 of wealth. We now sit George down in a room and offer him two investments. The first offers a sure-fire payoff of $1,000; the second offers some probability *p* of receiving $1,200 and a probability of *1 – p* of receiving $900. George's task is to decide what probability *p* he would need to make him indifferent between the two investments. Suppose he decides that he would only accept the risky investment if there was a 66.5% chance of the high payoff. Then, if he is maximizing utility, we know that the two investments must have the same expected utility, so that:

100 = .665 x utility of $1,200 + .335 x 80

Utility of $1,200 = 110.06

Similarly, we could ask George to consider an investment that paid either $900 or $1,140 and get him to state what odds of the higher payoff he would need to be indifferent between this risky investment and a sure-fire $1,000. That would allow us to calculate the utility that he places on wealth of $1,140. Suppose that the answer is that $1,140 provides a utility of 108.77. Now we can assign a utility to each of the payoffs from Portfolios A and B that George was contemplating earlier.

Payoff $	Utility
900	80
1,000	100
1,140	108.77
1,200	110.06

Portfolio A offers a 60% chance of a $1,200 payoff and a 40% chance of a $900 payoff. Its *expected* payoff is $1,080, but its expected utility to George is (.6 x 110.06) + (.4 x 80) = 98.04. Portfolio B offers a 50% chance of a $1,140 payoff and a 50% chance of a $1,000 payoff. Its *expected* payoff is $1,070, but its expected utility to George is (.5 x 108.77) + (.5 x 100) = 104.38. George derives a higher expected utility from portfolio B than from A.

Notice that the units in which we measure utility are arbitrary. The important thing is the *relative* difference in the utility that George derives from different payoffs. We could double each value or add 100 and we would reach the same predictions as to which portfolio George would choose.

Utility Functions

There is no need to confine ourselves to four possible payoffs. If George has the patience, we could ask him to make a series of choices between different hypothetical investments, and plot the utility that he derives from different levels of wealth. For example, we might find that George's choices are consistent with the utility function shown in Figure 1 below.

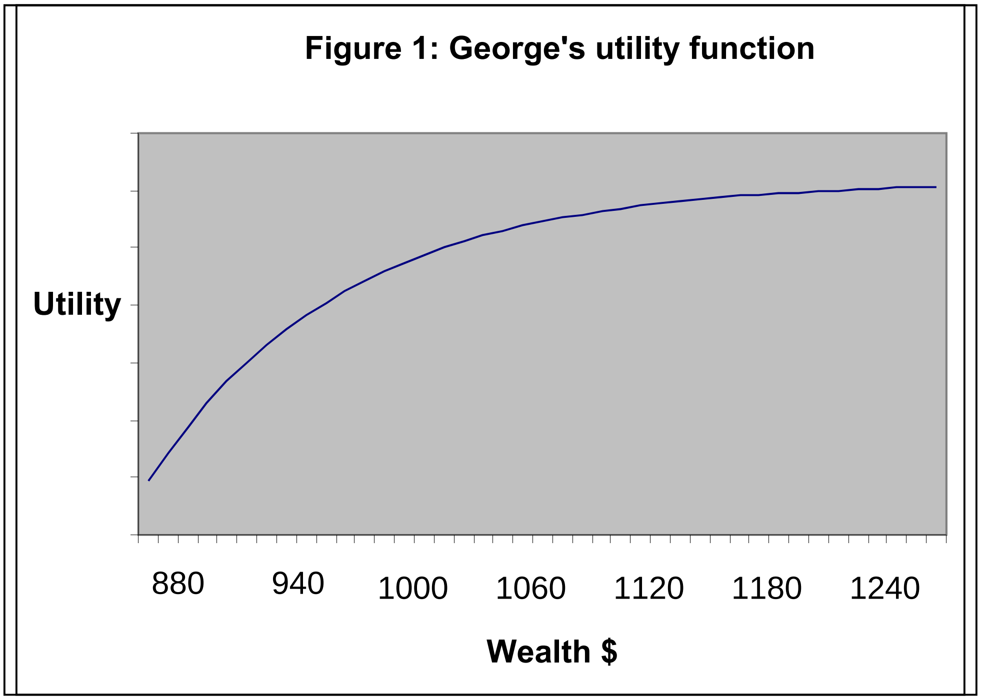

There are two things to notice about George's utility function. First, it slopes upward. In other words, George always prefers more wealth to less. Second, it is concave, so that a small reduction in wealth has a much greater effect on George's utility than a corresponding increase in wealth. In other words, George is risk averse.

Economists have suggested a number of possible forms for the utility function. For example, the plot in Figure 1 assumes that George's preferences for wealth of w can be described by a utility function of the form Utility = a + b x $e^{-\lambda w}$. One alternative is to assume that an investor's preferences are consistent with a quadratic function of the form Utility = Wealth – λWealth^2. Figure 2 plots preferences for an investor with a quadratic utility function. Notice that this plot has an important unsatisfactory feature. At high levels of wealth utility declines – the investor prefers less money to more.

Figure 2: A quadratic utility function

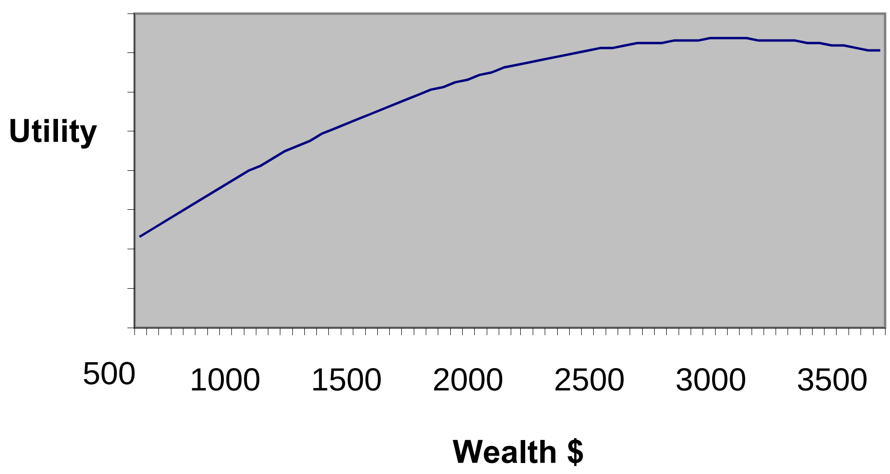

The Mean-Variance Assumptions

Mean-variance theory assumes that in choosing between two portfolios, an investor is concerned only with the mean outcome and the variance. This makes sense only in two special cases. One is that the portfolio returns are normally distributed and that the investor's utility function is exponential as in Figure 1. The other is that the distribution of returns can be of any form but that the investor's utility function is quadratic as in Figure 2. Neither assumption is wholly realistic. Stock returns are not normally distributed in the long run – you can't lose more than 100% but you might gain more than 100%. The alternative assumption is open to the criticism that it is unrealistic to assume that there is some level of wealth beyond which people would prefer less to more. Despite these drawbacks, mean-variance theory continues to play a central role in financial economics. It does so for two reasons. The notion that investors prefer higher expected return and dislike variance captures the essence of investor behavior. Second, if investor portfolio choice depends solely on the mean and variance, then it becomes possible to make a number of very important and practical predictions about optimal portfolios and the return that investors require for taking on risk.