What Practitioners Need to Know . . .

. . . About Utility

Mark Kritzman
Windham Capital Management

An important axiom of modern financial theory is that rational investors seek to maximize expected utility. Many financial analysts, however, find the concept of utility somewhat nebulous. This column discusses the origin of utility theory as well as its application within the context of financial analysis.

In his classic paper, "Exposition of a New Theory on the Measurement of Risk," first published in 1738, Daniel Bernoulli proposed the following: "the determination of the value of an item must not be based on its price, but rather on the utility it yields. The price of the item is dependent only on the thing itself and is equal for everyone; the utility, however, is dependent on the particular circumstances of the person making the estimate. Thus there is no doubt that a gain of one thousand ducats is more significant to a pauper than to a rich man though both gain the same amount."[1]

Bernoulli's insight that the utility of a gain depends on one's wealth may seem rather obvious and perhaps even pedestrian, yet it has profound implications for the theory of risk. Figure A helps us to visualize Bernoulli's notion of utility, which economists today call diminishing marginal utility.

The horizontal axis represents wealth, while the vertical axis represents utility. The relation between wealth and utility is measured by the curved line. Utility clearly increases with wealth, because the curve has a positive slope. The positive slope simply indicates that we prefer more

Figure A Diminishing Marginal Utility

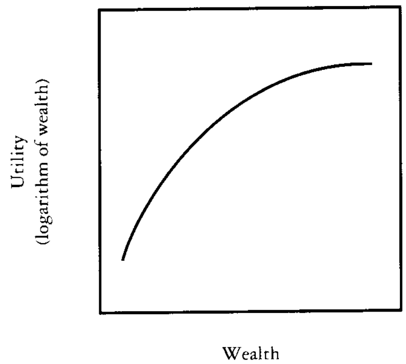

Figure B Changes in Utility versus Changes in Wealth

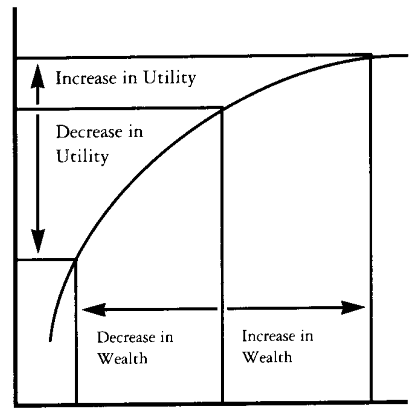

wealth to less wealth. This assumption seldom invites dispute.

As wealth increases, the increments to utility become progressively smaller, as the concave shape of the curve reveals. This concavity indicates that we derive less and less satisfaction with each subsequent unit of incremental wealth. Technically, Bernoulli's notion of utility implies that its first derivative with respect to wealth is positive, while its second derivative is negative.

A negative second derivative implies that we would experience greater *dis*utility from a decline in wealth than the utility we would derive from an equal increase in wealth. This tradeoff is apparent in Figure B, which shows the changes in utility associated with an equal increase and decrease in wealth.

According to Bernoulli, the precise change in utility associated with a change in wealth equals the logarithm of the sum of initial wealth plus the increment to wealth, divided by initial wealth. For example, the increase in utility associated with an increase in wealth from \$100 dollars to \$150 equals 0.405465, as follows:

$$0.405465 = \ln\left(\frac{100 + 50}{100}\right).$$

The next \$50 increment to wealth, however, yields a smaller increment to utility:

$$0.287682 = \ln\left(\frac{150 + 50}{150}\right)$$

Risk Aversion

From Bernoulli's assumption of diminishing marginal utility, it follows that those whose utility is defined by the logarithm of wealth will reject a "fair game"—that is, a game in which the expected outcome is the same for each participant. Bernoulli offered the following example. Two participants, each of whom has \$100 (Bernoulli used ducats), contribute \$50 as a stake in a game in which the chance of winning is determined by the toss of a coin. Each player therefore has a 50% probability of ending up

Table I Expected Utility of Participating in Fair Game

Payoff	Utility of Payoff	Probability	Probability-Weighted Payoff
50	3.91	50%	1.96
150	5.01	50%	2.50
Expected Utility			4.46

with $50 and a 50% probability of ending up with $150.

The expected *value* of this game is $100 (0.5 · 50 + 0.5 · 150). This is identical to the expected value of *not* playing, because each participant starts out with $100. The expected *utility* of this game, however, is found by adding the utility (logarithm) of the $50 payoff times its probability of occurring to the utility of the $150 payoff times its probability of occurring.

Table I shows the results. The utility of not participating in the game equals the logarithm of $100 dollars, which is 4.61. The utility of *not* participating is higher than the utility of participating. Thus a player whose utility with respect to wealth is defined as the logarithm of wealth will reject the game, even though it is fair.

In essence, Bernoulli's concept of utility implies that we prefer a certain prospect to an uncertain prospect of equal value. That is to say, we are risk averse. In fact, Bernoulli interpreted risk aversion as "nature's admonition to avoid the dice."[2]

Certainty Equivalent

With simple algebra, we can extend Bernoulli's insight a little further and determine how much value we would be willing to subtract from a certain prospect before we would select a risky prospect. The value of the certain prospect that yields the same utility as the expected utility of an uncertain prospect is called the certainty equivalent.

We have already seen that the expected utility of a risky prospect is computed as the logarithm of the favorable outcome times its probability of occurring plus the logarithm of the unfavorable outcome times its probability of occurring (which would equal one minus the probability of occurrence of the favorable outcome). We can set this formula equal to the logarithm of the certain prospect and solve for the value of the certain prospect, as follows:

$$\ln(C) = \ln(F) \cdot p + \ln(U) \cdot (1 - p)$$

$$C = e^{[\ln(F) \cdot p + \ln(U) \cdot (1 - p)]}$$

where:

ln = natural logarithm,
C = certain payoff,
F = payoff from favorable outcome,
U = payoff from unfavorable outcome,
p = probability of occurrence for favorable outcome and
e = 2.71828.

In Bernoulli's example, in which each participant has a 50% chance of receiving $150 and a 50% chance of receiving $50, the certainty equivalent is $86.60:

$$C = e^{[\ln(150) \cdot 0.5 + \ln(50) \cdot (1 - 0.5)]}$$

$$86.60 = e^{(2.5053 + 1.9560)}.$$

The certainty equivalent of $86.60 implies that we would be be indifferent between a certain $86.60 payoff and a risky prospect with an equal probability of paying $150 or $50. It follows that we would select the risky prospect only if it offered a higher expected value than the certain prospect. The difference in expected value that would induce us to choose the risky prospect is called the required risk premium.

We have thus far assumed that our initial wealth equals only $100. Assume, instead, that we were starting with initial wealth of $1,000 dollars, so that we would receive $1,050 given a favorable outcome and $950 given an unfavorable outcome. The certainty equivalent would then equal $998.75. Whereas a person with only $100 would demand a risk premium of more than 13% ([100 − 86.60]/100), a person with $1,000 would demand a risk premium of only one-eighth of 1% ([1,000 − 998.75]/1000). The wealthier person is still risk averse, but he is not nearly as disinclined as the poorer one to incur the risk of losing $50, because this amount represents but a small fraction of his wealth. Of course, if the potential gain or loss from the risky prospect was $500, he too would prefer a 13% risk premium.

It is easy to see how this framework also enables us to determine how much we should pay to insure against various risks. Suppose, for example, that we have $100,000 of savings and we inherit a family heirloom valued at $10,000, which is to be mailed to us. Suppose further that there is a 1% chance that the heirloom will be lost in the mail. Based on Bernoulli's model of risk aversion, we should be willing to pay up to $104.79 to insure this heirloom.

The total payoff, should the heirloom arrive safely, equals $110,000, which yields utility of 11.6082. If the heirloom is lost in the mail, the total payoff equals $100,000, which yields utility of 11.5129. Thus the expected utility of this risky prospect equals 11.6073 (11.6082 · 0.99 + 11.5129 · 0.01). This expected utility corresponds to a certainty equivalent of $109,895.21. Thus we should be willing to pay up to

Table II Risk Aversion

	Absolute	Relative
Decreasing	Increase Risky Amount	Increase Risky Percentage
Constant	Maintain Risky Amount	Maintain Risky Percentage
Increasing	Decrease Risky Amount	Decrease Risky Percentage

$104.79 (110,000 − 109,895.21) to insure the heirloom.

The Rhetoric of Risk Preference

Bernoulli's view of utility is certainly plausible, but we should not conclude that it describes everyone's attitude toward risk. Economists have generalized Bernoulli's insights into a comprehensive theory of risk preference accompanied by the usual classifications and rhetoric. For example, economists distinguish between those who are risk averse, those who are risk neutral and those who seek risk. A risk-averse person will reject a fair game, while a risk-neutral person will be indifferent to a fair game and a risk seeker will select a fair game. Economists also distinguish between the absolute amount of one's wealth that is exposed to risk versus the proportion of one's wealth that is exposed to risk, and whether this amount decreases, remains constant or increases as wealth increases.

Decreasing absolute risk aversion indicates that the amount of wealth one is willing to expose to risk increases as one's wealth increases. Constant absolute risk aversion implies that the amount of wealth exposed to risk remains unchanged as wealth increases. Increasing absolute risk aversion means that absolute risk exposure decreases as wealth increases.

Relative risk aversion refers to changes in the percentage of one's wealth exposed to risk as wealth increases. Decreasing relative risk aversion implies that the percentage of wealth exposed to risk increases as wealth increases. With constant relative risk aversion, the percentage of wealth exposed to risk does not change as wealth increases. Increasing relative risk aversion implies that percentage risk exposure decreases as wealth increases. Table II summarizes these relationships.

Indifference Curves

In many applications, it is convenient to model utility as a function of expected return and risk as measured by the standard deviation of returns. Because most investors are indeed risk averse, utility is usually depicted as a positive function of expected return and a negative function of risk:[3]

$$U = E(r) - \lambda \cdot \sigma^2$$

where:

U = utility,
$E(r)$ = expected return,
σ = standard deviation of returns and
λ = risk-aversion coefficient.

The risk-aversion coefficient has no economic meaning in and of itself. It is merely an index of our aversion toward risk. Of course, if we are risk averse, the coefficient must be positive. And the higher the value of the coefficient, the more risk averse we would be.

Suppose, for example, that our risk-aversion coefficient equals 5. An asset with an expected return of 8% and a standard deviation of 10% would yield utility of 3.0%. Another asset may have an expected return of 10% and a standard deviation of 12%. The utility from this asset, given our aversion toward risk, would equal 2.8%. The return of the second asset is not high enough to compensate for its higher risk, even though its expected return and risk are both 2% higher than those of the first asset. Given a risk-aversion coefficient equal to 5, we would prefer the less risky asset.

If we were less risk averse, with a coefficient of 3, say, the first asset would yield utility of 5.0%, while the second asset would yield utility of 5.7%. In this case, the incremental expected return of the second asset is sufficient to counter its higher risk. Table III summarizes these results.

It is possible to identify combinations of expected return and standard deviation that yield the same level of utility for a particular risk-aversion coefficient. Table IV shows several combinations based on a risk-aversion coefficient of 4. It indicates that we are more inclined to incur risk in order to increase expected return at low levels of expected return than we are at higher levels of expected return. For example, if we are starting from an expected return of 5%, we are willing to accept five additional units of risk

Table III The Risk-Aversion Coefficient

	Expected Return	Standard Deviation	Utility	Preference
Coefficient = 5				
Asset 1	8.0%	10.0%	3.0%	←
Asset 2	10.0%	12.0%	2.8%	
Coefficient = 3				
Asset 1	8.0%	10.0%	5.0%	
Asset 2	10.0%	12.0%	5.7%	←

Table IV Risk-Return Combinations with Equal Utility (risk aversion = 4)

Expected Return	*Standard Deviation*	*Utility*
5.00	0.00	5.00
6.00	5.00	5.00
7.00	7.07	5.00
8.00	8.66	5.00
9.00	10.00	5.00
10.00	11.18	5.00
11.00	12.25	5.00
12.00	13.23	5.00
13.00	14.14	5.00
14.00	15.00	5.00
15.00	15.81	5.00
16.00	16.58	5.00
17.00	17.32	5.00
18.00	18.03	5.00
19.00	18.71	5.00
20.00	19.56	5.00

in order to increase our expected return by one unit. If, however, we are starting from an expected return of 15%, we are willing to accept only 0.77 units of incremental risk in order to gain one additional unit of expected return.

Tracing a curve through all the combinations of expected return and risk with equal utility creates an indifference curve. Figure C shows three hypothetical indifference curves. Indifference curves that are closer to the upper left corner yield more utility and thus are more desirable. Given a particular indifference curve, however, all the points along that curve yield the same utility.

Figure C Indifference Curves

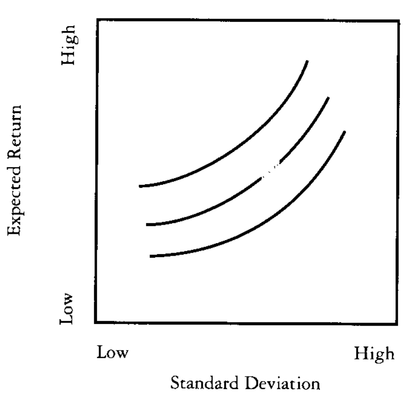

Figure D The Optimal Portfolio

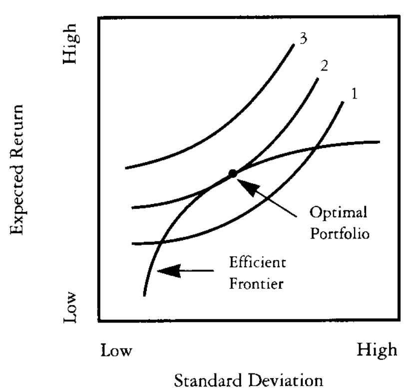

If we were to combine a set of risky assets efficiently so that, for any given level of expected return, we minimized risk, a continuum of such combinations would form a concave curve in dimensions of expected return and risk. Figure D shows such a curve, which is called the efficient frontier, along with three indifference curves.[4] The point of tangency between the efficient frontier and indifference curve 2 represents the precise combination of expected return and standard deviation that maximizes utility. It matches our preference for incurring risk in order to raise expected return with the best available tradeoff of risk and return from the capital markets.

Clearly, we would prefer a combination of expected return and risk located along indifference curve 3, but indifference curve 3 is located in a region that is unobtainable. Indifference curve 1, on the other hand, is undesirable because it is dominated by many of the combinations along the efficient frontier. This illustration shows how we can employ utility theory to identify a portfolio of assets that is optimal given our particular attitude toward risk.

The concept of utility is critical to the theory of choice under uncertainty. While it is reasonable to accept the basic premises of utility theory, such as the notion that more wealth is preferred to less wealth and that investors are typically risk averse, it is important to recognize that concise mathematical models of utility may not always reflect the full range of investor attitudes and idiosyncracies. We should be sensitive to the fact that some descriptions of utility are put forth primarily for tractability or expository convenience.[5]

Footnotes

1. *D. Bernoulli, "Exposition of a New Theory on the Measurement of Risk,"* Econometrica, *January 1954 (translation from 1738 version). Daniel Bernoulli is one of several celebrated Bernoulli mathematicians. Daniel's father, Johann, made important contributions to calculus, although much of his work was published by the Marquis de l'Hospital. Johann was also the mentor of the famous prodigy Leonhard Euler. The most renowned Bernoulli was Daniel's uncle and Johann's older brother Jakob. Jakob Bernoulli is known primarily for his contributions to the theory of probability. Finally, Daniel's cousin, Nicolas Bernoulli, was a distinguished mathematician who proposed the famous St. Petersburg Paradox, for which Daniel offered a solution in his classic risk measurement paper.*
2. Ibid., *p. 29.*
3. *The capital markets offer compelling evidence of risk aversion. Historically, investors have priced assets to extract a premium for bearing risk; to wit, annually from 1926 through 1991, stocks have returned 5.6% more than government bonds, and government bonds have returned 1.1% more than Treasury bills.*
4. *For a more detailed discussion of the efficient frontier, see M. Kritzman, "What Practitioners Need to Know . . . About The Nobel Prize,"* Financial Analysts Journal, *January/February 1991.*
5. *For a basic review of utility theory, see E. Elton and M. Gruber,* Modern Portfolio Theory and Investment Analysis, *third edition (New York: John Wiley & Sons, 1987), pp. 179–203. For a more technical discussion of utility theory, see C. Huang and R. Litzenberger,* Foundations for Financial Economics *(New York: North-Holland, 1988), pp. 1–37.*