

An Introduction to Utility Theory

John Norstad

j.norstad@mac.com

<http://homepage.mac.com/j.norstad>

March 29, 1999

Updated: January 22, 2005

Abstract

A gentle but reasonably rigorous introduction to utility theory.

Utility functions give us a way to measure investor's preferences for wealth and the amount of risk they are willing to undertake in the hope of attaining greater wealth. This makes it possible to develop a theory of portfolio optimization. Thus utility theory lies at the heart of modern portfolio theory.

We develop the basic concepts of the theory through a series of simple examples. We discuss non-satiation, risk aversion, the principle of expected utility maximization, fair bets, certainty equivalents, portfolio optimization, coefficients of risk aversion, iso-elasticity, relative risk aversion, and absolute risk aversion.

Our examples of possible investments are deliberately over-simplified for the sake of exposition. While they are much too simple to be directly relevant for real-life applications, they lay the foundation upon which the more complicated relevant theories are developed.

This paper was inspired by the first few pages of chapter 2 of Robert Merton's book [1]. It is a student's clumsy attempt to fill in the gaps in Merton's presentation of this material. Merton assumes basic utility theory as a given. For your ignorant author it was anything but a given. Our notation and terminology are largely that of Merton's.

Contents

1	Introduction	2
2	Example 1 – A Fair Game	3
3	Positive Affine Transformations	6
4	Example 2 – An All-or-Nothing Investment	7
5	Example 3 – Optimizing a Portfolio	10
6	The Logarithmic Utility Function	15
7	The Iso-Elastic Utility Functions	16
8	The Negative Exponential Utility Function	18
9	Risk Aversion Functions	20
10	Example 4 – Variable Risk Aversion	22
11	Utility and Real-Life Investors	25

List of Figures

1	Example 1 – A Fair Game	4
2	Example 2 – An All-or-Nothing Investment, $U(w) = -10^6/w^3$. .	8
3	Example 2 – An All-or-Nothing Investment, $U(w) = -10^{10}/w^5$.	9
4	Example 3 – Utility Hill for $\lambda = -3$	11
5	Example 3 – Utility Hill for $\lambda = -5$	12
6	Example 3 – Optimizing a Portfolio	14
7	Example 4 – Variable Risk Aversion	23

1 Introduction

A *utility function* is a twice-differentiable function of wealth $U(w)$ defined for $w > 0$ which has the properties of *non-satiation* (the first derivative $U'(w) > 0$) and *risk aversion* (the second derivative $U''(w) < 0$).¹

A utility function measures an investor's relative preference for different levels of total wealth.

The non-satiation property states that utility increases with wealth, i.e., that more wealth is preferred to less wealth, and that the investor is never satiated – he never has so much wealth that getting more would not be at least a little bit desirable.

The risk aversion property states that the utility function is concave or, in other words, that the marginal utility of wealth decreases as wealth increases.

To see why utility functions are concave, consider the extra (marginal) utility obtained by the acquisition of one additional dollar. For someone who only has one dollar to start with, obtaining one more dollar is quite important. For someone who already has a million dollars, obtaining one more dollar is nearly meaningless. In general, the increase in utility caused by the acquisition of an additional dollar decreases as wealth increases.

It may not be obvious what this concavity of utility functions or “decreasing marginal utility of wealth” has to do with the notion of “risk aversion.” This should become clear in the examples.

Different investors can and will have different utility functions, but we assume that any such utility function satisfies the two critical properties of non-satiation and risk aversion.

The *principle of expected utility maximization* states that a rational investor, when faced with a choice among a set of competing feasible investment alternatives, acts to select an investment which maximizes his expected utility of wealth.

Expressed more formally, for each investment I in a set of competing feasible investment alternatives F , let $X(I)$ be the random variable giving the ending value of the investment for the time period in question. Then a rational investor with utility function U faces the optimization problem of finding an investment $I_{opt} \in F$ for which:

$$E(U(X(I_{opt}))) = \max_{I \in F} E(U(X(I)))$$

¹It some contexts it is necessary to consider utility functions which are defined for zero and negative values of wealth. For simplicity we ignore this complication.

2 Example 1 – A Fair Game

As a first example, consider an investor with a square root utility function:

$$U(w) = \sqrt{w} = w^{0.5}$$

Note that :

$$\begin{aligned} U'(w) &= 0.5w^{-0.5} > 0 \\ U''(w) &= -0.25w^{-1.5} < 0 \end{aligned}$$

So this is a legitimate utility function.

We assume that the investor's current wealth is \$5.

To make things simple, we assume that there is only one investment available. In this investment a fair coin is flipped. If it comes up heads, the investor wins (receives) \$4, increasing his wealth to \$9. If it comes up tails, the investor loses (must pay) \$4, decreasing his wealth to \$1. Note that the expected gain is $0.5 \times (\$4) + 0.5 \times (-\$4) = \$0$. This is called a “fair game.”

This game may seem more like a “bet” than an “investment.” To see why it's an investment, consider an investment which costs \$5 (all of our investor's current wealth) and which has two possible future values: \$1 in the bad case and \$9 in the good case. This investment is clearly exactly the same as the coin-flipping game.

Note that we have chosen a very volatile investment for our example. In the bad case, the rate of return is -80%. In the good case, the rate of return is +80%. Note that the expected return is 0%, as is the case in all fair games/bets/investments.

We assume that our investor has only two choices (the set of feasible investment alternatives has only two elements). The investor can either play the game or not play the game (do nothing). Which alternative does the investor choose if he follows the principle of expected utility maximization?

Figure 1 shows our investor's current wealth and utility, the wealth and utility of the two possible outcomes in the fair game, and the expected outcome and the expected utility of the outcome in the fair game.

If the investor refuses to play the game and keeps his \$5, he ends up with the same \$5, for an expected utility of $\sqrt{5} = 2.24$. If he plays the game, the expected outcome is the same \$5, but the expected utility of the outcome is only $0.5 \times 1 + 0.5 \times 3 = 2$. Because he acts to maximize expected utility, and because 2.24 is greater than 2, he refuses to play the game.

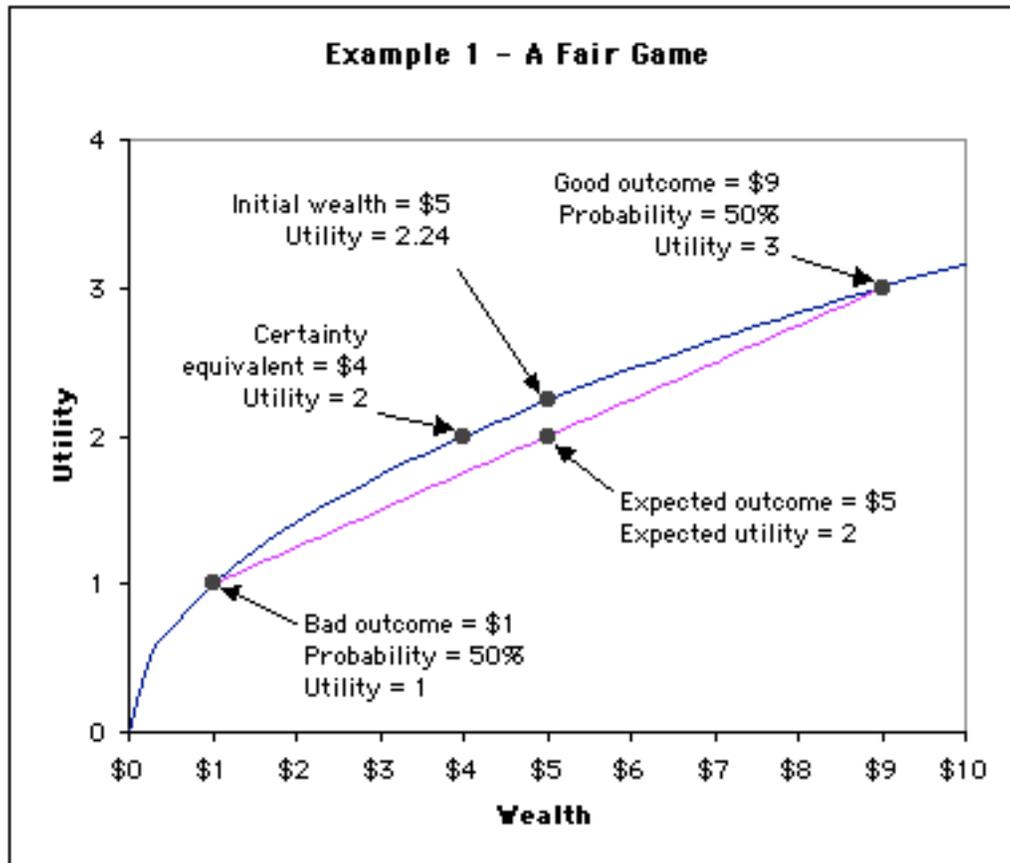


Figure 1: Example 1 – A Fair Game

In general, a risk-averse investor will always refuse to play a fair game where the expected return is 0%. If the expected return is greater than 0%, the investor may or may not choose to play the game, depending on his utility function and initial wealth.

For example, if the probability of the good outcome in our example was 75% instead of 50%, the expected outcome would be \$7, the expected gain would be \$2, the expected return would be 40%, and the expected utility would be 2.5. Because 2.5 is greater than 2.24, the investor would be willing to make the investment. The expected return of 40% is a “risk premium” which compensates him for undertaking the risk of the investment.

Another way of looking at this property of risk aversion is that investors attach greater weight to losses than they do to gains of equal magnitude. In the example above, the loss of \$4 is a decrease in utility of 1.24, while the gain of \$4 is an

increase in utility of only 0.76.

Similarly, for a risk-averse investor, a loss of $2x$ is more than twice as bad as a loss of $1x$, and a gain of $2x$ is less than twice as good as a gain of $1x$. In the example above, with an initial wealth of \$5, a loss of \$1 is a decrease in utility of 0.24, and a loss of \$2 is a decrease in utility of 0.50, more than twice 0.24. A gain of \$1 is an increase in utility of 0.21, and a gain of \$2 is an increase in utility of 0.41, less than twice 0.21.

In our example, the expected utility of the outcome is 2. The wealth value which has the same utility is \$4 (2 squared). This value \$4 is called the *certainty equivalent*. If the initial wealth is less than \$4, an investor with a square root utility function would choose to play a game where the outcome is an ending wealth of \$1 with probability 50% (a loss of less than \$3) or an ending wealth of \$9 with probability 50% (a gain of more than \$5). Put another way, with a current wealth of \$4, our investor is willing to risk the loss of about 75% of his current wealth in exchange for an equal chance at increasing his wealth by about 125%, but he's not willing to risk any more than this.

In general, the certainty equivalent for an investment whose outcome is given by a random variable X is:

$$\text{Certainty equivalent} = c = U^{-1}(E(U(X)))$$

$$U(c) = E(U(X))$$

If an investor with utility function U has current wealth less than c , he will consider the investment attractive (although some other investment may be even more attractive). If his current wealth is greater than c , he will consider the investment unattractive, because doing nothing has greater expected utility than the investment. If his current wealth is exactly c , he will be indifferent between undertaking the investment and doing nothing.

Note that because U is an increasing function, maximizing expected utility is equivalent to maximizing the certainty equivalent.

The certainty equivalent is always less than the expected value of the investment. In our example, the certainty equivalent is \$4, while the expected value (expected outcome) is \$5.

3 Positive Affine Transformations

Utility functions are used to compare investments to each other. For this reason, we can scale a utility function by multiplying it by any positive constant and/or translate it by adding any other constant (positive or negative). This kind of transformation is called a *positive affine* transformation.

For example, with the square root utility function we used above, we could have used any of the following functions instead:

$$\begin{aligned} &100\sqrt{w} \\ &50\sqrt{w} + 83 \\ &\sqrt{w} - 413 \\ &\frac{3\sqrt{w} + 10}{8} \\ &\text{etc.} \end{aligned}$$

The specific numbers appearing as the utility function values on our graphs and in our calculations would be different, but the graphs would all look the same when scaled and translated appropriately, and all our results would be the same.

It is easy to see why this is true in general. Suppose we have constants $a > 0$ and b and a utility function U . Define another utility function V :

$$V(w) = aU(w) + b$$

Note that:

$$\begin{aligned} V'(w) &= aU'(w) > 0 && \text{because } a > 0 \text{ and } U'(w) > 0 \\ V''(w) &= aU''(w) < 0 && \text{because } a > 0 \text{ and } U''(w) < 0 \end{aligned}$$

so V is a valid utility function.

Consider an investment I with outcome given by a random variable X . We can easily see that the certainty equivalent for I is the same under utility function V as it is under utility function U . Let c be the certainty equivalent under U . Then:

$$\begin{aligned} c &= U^{-1}(E(U(X))) \\ U(c) &= E(U(X)) \\ V(c) &= aU(c) + b = aE(U(X)) + b = E(V(X)) \\ c &= V^{-1}(E(V(X))) \end{aligned}$$

When we talk about utility functions we often say that two functions are the “same” when they differ only by a positive affine transformation.

4 Example 2 – An All-or-Nothing Investment

For our second example we assume that the investor's current wealth is \$100 and we begin by using the following utility function:

$$\begin{aligned} U(w) &= \frac{-1,000,000}{w^3} = -1,000,000w^{-3} \\ U'(w) &= 3,000,000w^{-4} > 0 \\ U''(w) &= -12,000,000w^{-5} < 0 \end{aligned}$$

Our utility function is the same as $-1/w^3$. We use the scale factor 1,000,000 to make the function values easier to read in the wealth neighborhood of \$100 which we are investigating. Without the scale factor, the numbers are very small, messy to write, and not easy to interpret at a glance.

As in the first example, we assume that our investor has only one alternative to doing nothing. He may use his entire wealth of \$100 to purchase an investment which returns -10% with probability 50% and +20% with probability 50%. We will continue to use this hypothetical investment as a running example through the rest of this paper.

Note that the expected return on this investment is +5% and the standard deviation of the returns is 15%. This is similar to many common real-life financial investments, except that in real life there are many more than only two possible outcomes.²

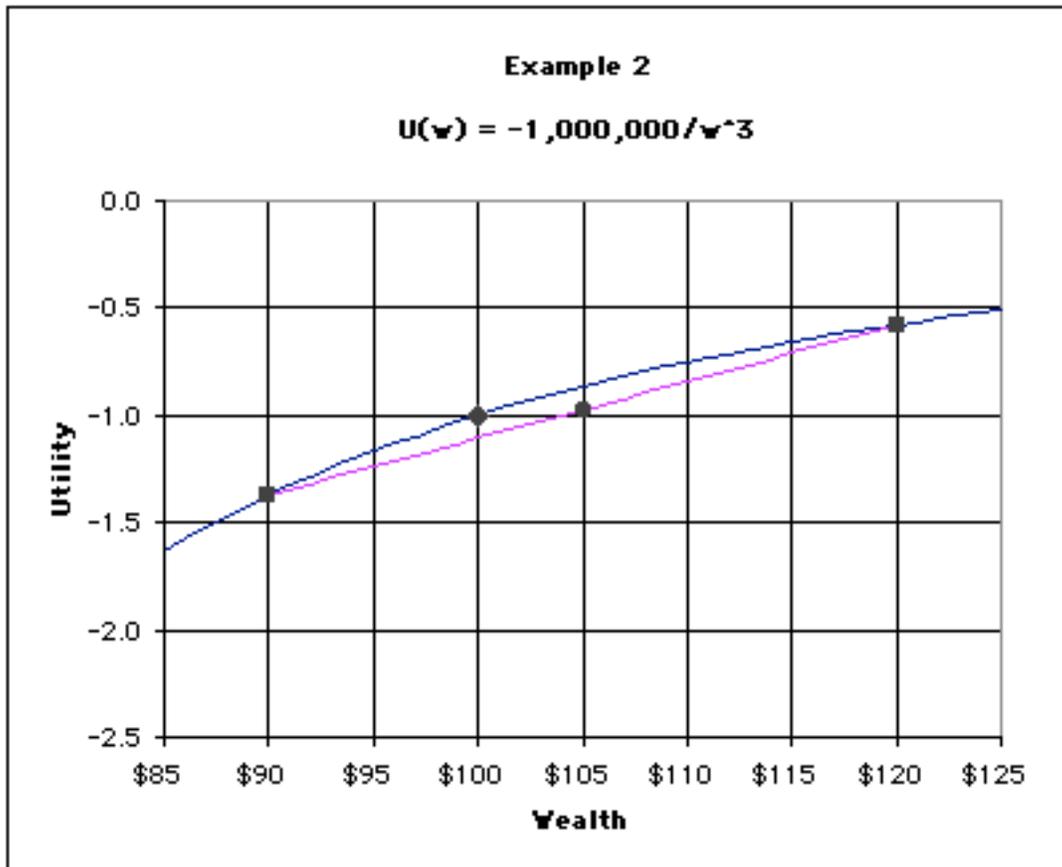
We continue to assume that there are only two possible outcomes to make the math easier for the sake of exposition.

Figure 2 shows a graph similar to the one in our first example.

The expected utility of the investment is -0.98, which is larger than the utility of doing nothing, which is -1.00. Thus, in this case the investor chooses to make the investment.

The decision looks like a close call in this example. The expected utility of the investment is only slightly larger than that of doing nothing. This leads us to wonder what might happen if we change the exponent in the utility function.

²We do not consider issues of time or the time value of money. Our risk-free "do nothing" alternative earns no interest. In the real world we would earn a risk-free interest rate. Thus our 5% return should be compared to excess returns above the risk-free interest rate in the real world.

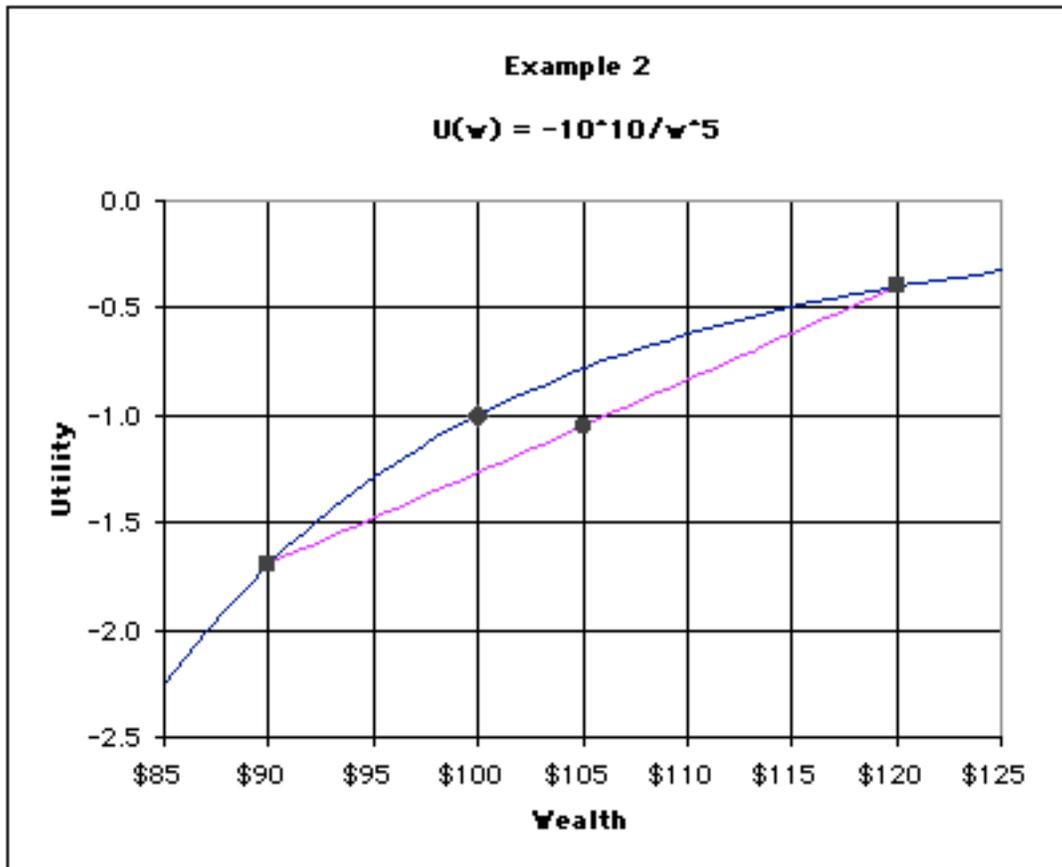


Wealth	Utility	
\$90	-1.37	Bad outcome
\$100	-1.00	Current wealth
\$105	-0.98	Expected outcome
\$120	-0.58	Good outcome

Figure 2: Example 2 – An All-or-Nothing Investment, $U(w) = -10^6/w^3$

The graph in Figure 3 shows how the situation changes if we use an exponent of -5 instead of -3 in the utility function.

$$\begin{aligned}
 U(w) &= \frac{-10^{10}}{w^5} = -10^{10}w^{-5} \\
 U'(w) &= 5 \times 10^{10}w^{-6} > 0 \\
 U''(w) &= -30 \times 10^{10}w^{-7} < 0
 \end{aligned}$$



Wealth	Utility	Description
\$90	-1.69	Bad outcome
\$100	-1.00	Current wealth
\$105	-1.05	Expected outcome
\$120	-0.40	Good outcome

Figure 3: Example 2 – An All-or-Nothing Investment, $U(w) = -10^{10}/w^5$

In this case, the expected utility of the investment is -1.05, which is smaller than the utility of doing nothing, which is -1.00. Thus, in this case the investor chooses not to make the investment.

This example shows that an investor with a utility function of $-1/w^5$ is somewhat more risk-averse than is an investor with a utility function of $-1/w^3$.

5 Example 3 – Optimizing a Portfolio

In the first two examples we assumed that the investor had only two options: do nothing or invest all of his money in a risky asset.

In this example we use the same kind of utility functions as in examples 1 and 2 and the same investment as in example 2. This time, however, we permit investing any desired portion of the investor's total wealth of \$100 in the risky asset.

The investor may choose to do nothing, invest everything in the risky asset, or do nothing with part of his money and invest the rest.

We also generalize the utility function in example 2 to permit exponents other than -3 and -5. For reasons which will become clear later we use the following parameterized form of these utility functions:

For any $\lambda < 1, \lambda \neq 0$:

$$\begin{aligned} U_\lambda(w) &= \frac{w^\lambda - 1}{\lambda} \\ U'_\lambda(w) &= w^{\lambda-1} > 0 \\ U''_\lambda(w) &= (\lambda - 1)w^{\lambda-2} < 0 \end{aligned}$$

Note that we must have $\lambda < 1$ to guarantee $U''_\lambda < 0$. If $\lambda = 1$ we get a risk-neutral utility function. If $\lambda > 1$ we get a risk-loving utility function. We assume risk aversion and do not investigate these alternatives here.

Consider the utility functions we used in examples 1 and 2: $\sqrt{w} = w^{0.5}$, $-w^{-3}$, and $-w^{-5}$. These functions are the same as $U_{0.5}$, U_{-3} , and U_{-5} .

Our investor has a current wealth of \$100 and may choose to invest any part of it in the risky asset. Let:

$$\alpha = \text{the amount invested in the risky asset}$$

Then $100 - \alpha$ is the amount with which the investor does nothing. The two possible outcomes are:

$$\begin{aligned} \text{Bad outcome: } w &= 0.9\alpha + (100 - \alpha) = 100 - 0.1\alpha \\ \text{Good outcome: } w &= 1.2\alpha + (100 - \alpha) = 100 + 0.2\alpha \end{aligned}$$

The expected utility of the outcome is:

$$\begin{aligned} f(\alpha) &= 0.5U(100 - 0.1\alpha) + 0.5U(100 + 0.2\alpha) \\ &= 0.5 \frac{(100 - 0.1\alpha)^\lambda - 1}{\lambda} + 0.5 \frac{(100 + 0.2\alpha)^\lambda - 1}{\lambda} \\ &= \frac{0.5}{\lambda} [(100 - 0.1\alpha)^\lambda + (100 + 0.2\alpha)^\lambda - 2] \end{aligned}$$

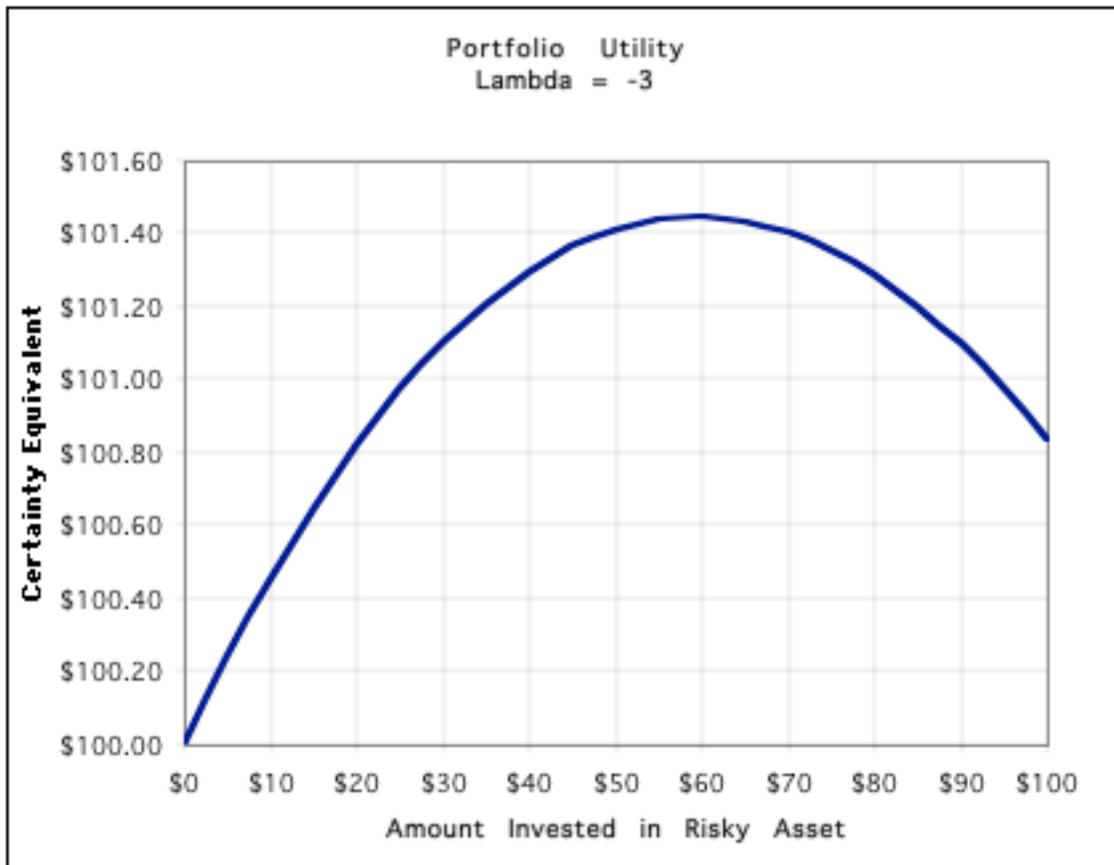


Figure 4: Example 3 – Utility Hill for $\lambda = -3$

Figure 4 displays the certainty equivalent of this function $U^{-1}(f(\alpha))$ for $\lambda = -3$. Recall that maximizing expected utility is equivalent to maximizing the certainty equivalent. Using certainty equivalents in optimization problems like this one is often more natural and intuitive than working directly with the expected utility values.

As we saw in example 2, given an all-or-nothing choice, our investor would chose to invest his \$100 in the risky asset. Given the opportunity to invest an arbitrary portion of his total wealth, however, we see that neither extreme choice is optimal. In this case, with $\lambda = -3$, the optimal amount to invest in the risky asset is about \$59. The investor does nothing with the remaining \$41.

Figure 5 shows the corresponding graph for $\lambda = -5$.

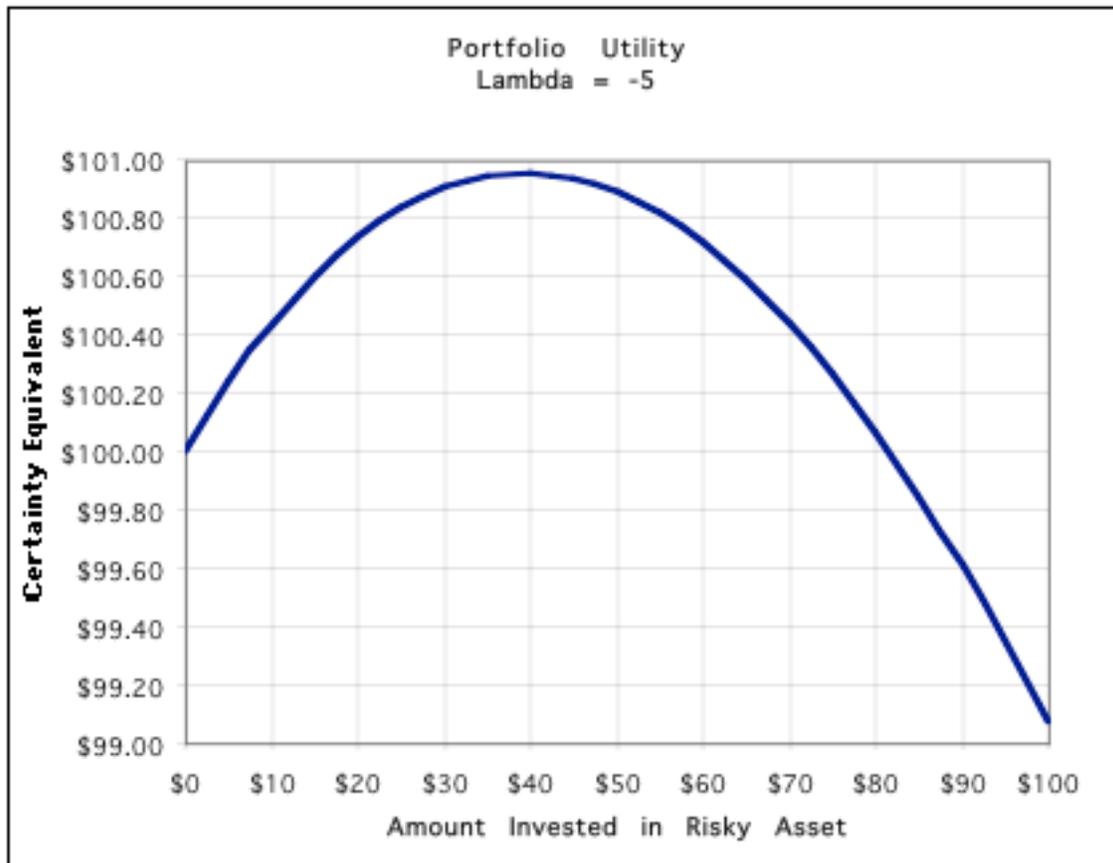


Figure 5: Example 3 – Utility Hill for $\lambda = -5$

In this case, the optimal portfolio for our somewhat more risk-averse investor is to invest \$39 in the risky asset and do nothing with the remaining \$61.

The curves in these graphs are called “utility hills.” In general, when we mix together multiple possible risky assets, the curves becomes surfaces in two or more dimensions.

The general asset allocation portfolio optimization problem is to climb the hill to find the particular portfolio at its peak.

In the very simple kind of portfolio we’re examining in this example, we can easily use a bit of calculus and algebra to find the exact optimal portfolio.

The principle of expected utility maximization tells us that we need to find the value of α for which $f(\alpha)$ attains its maximum value. To do this we take the derivative of f , set it equal to 0, and solve for α :

$$\begin{aligned} f'(\alpha) &= \frac{0.5}{\lambda} [-0.1\lambda(100 - 0.1\alpha)^{\lambda-1} + 0.2\lambda(100 + 0.2\alpha)^{\lambda-1}] \\ &= 0.1(100 + 0.2\alpha)^{\lambda-1} - 0.05(100 - 0.1\alpha)^{\lambda-1} = 0 \end{aligned}$$

Rearranging and simplifying we get:

$$\begin{aligned} 0.1(100 + 0.2\alpha)^{\lambda-1} &= 0.05(100 - 0.1\alpha)^{\lambda-1} \\ 2 &= \left(\frac{100 - 0.1\alpha}{100 + 0.2\alpha} \right)^{\lambda-1} = \left(\frac{100 + 0.2\alpha}{100 - 0.1\alpha} \right)^{1-\lambda} \end{aligned}$$

$\lambda < 1$, so the exponent $1 - \lambda$ in this last equation is greater than 0. This number is called the *coefficient of risk aversion* and is often denoted by the variable A :

$$A = 1 - \lambda = \text{coefficient of risk aversion}$$

As we saw in example 2, as λ decreases, investors become more risk-averse. Thus, as A increases, investors become more risk-averse.

For the two examples we saw above, in the case $\lambda = -3$ we have $A = 4$, and for $\lambda = -5$ we have $A = 6$. The investor with $A = 6$ is more risk-averse than is the investor with $A = 4$.

We now rewrite the last equation above using our new coefficient of risk aversion and solve the result for α :

$$\begin{aligned} \left(\frac{100 + 0.2\alpha}{100 - 0.1\alpha} \right)^A &= 2 \\ \frac{100 + 0.2\alpha}{100 - 0.1\alpha} &= 2^{1/A} \\ 100 + 0.2\alpha &= 100 \times 2^{1/A} - 0.1 \times 2^{1/A} \alpha \\ (0.2 + 0.1 \times 2^{1/A}) \alpha &= 100(2^{1/A} - 1) \\ \alpha &= \frac{100(2^{1/A} - 1)}{0.2 + 0.1 \times 2^{1/A}} \end{aligned}$$

We graph this function for the coefficient of risk aversion A ranging from 2 to 20 in Figure 6.

As expected, as risk aversion increases, the portion of the optimal portfolio which is invested in the risky asset gets smaller.

Something interesting happens at the left edge of this graph. For an investor with a low coefficient of risk aversion $A = 2$, the optimal amount to invest in the risky asset is about \$121. This is \$21 more than our investor's total wealth!

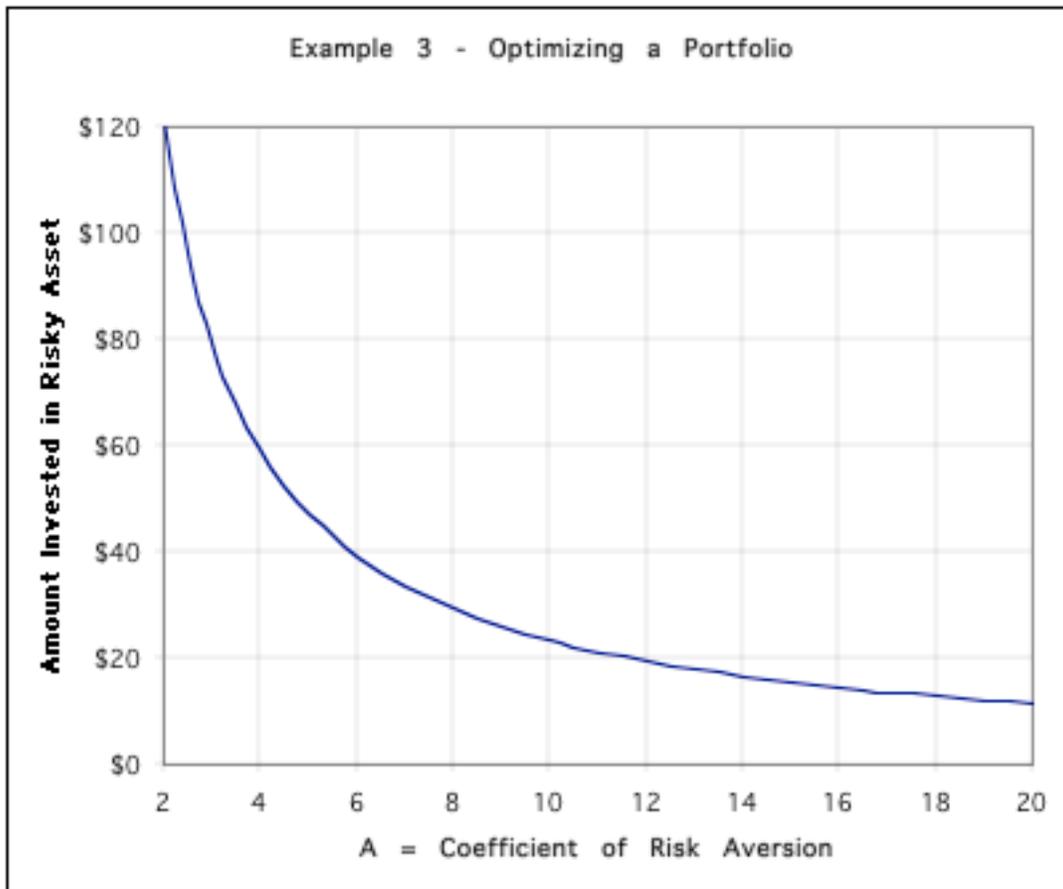


Figure 6: Example 3 – Optimizing a Portfolio

Suppose the investor is able to borrow an extra \$21 from someone without having to pay any interest on the loan. If this is the case, the optimal portfolio for our investor is to borrow the \$21, put it together with his \$100, and invest the resulting \$121 in the risky asset. For this investor with $A = 2$, however, it would not be optimal to borrow any more than \$21.

Borrowing money to help finance a risky investment is called “leverage.”

If our investor is unable to borrow money, his optimal portfolio is to invest his entire current wealth in the risky asset.

6 The Logarithmic Utility Function

In example 3 we looked at the class of utility functions:

$$U_\lambda(w) = \frac{w^\lambda - 1}{\lambda} \quad \lambda < 1 \text{ and } \lambda \neq 0$$

There is a conspicuous hole in this collection at $\lambda = 0$, corresponding to $A = 1$.

Fortunately, this hole is quite easily and elegantly filled by taking the limit of U_λ as $\lambda \rightarrow 0$ using L'Hôpital's rule:

$$\begin{aligned} \lim_{\lambda \rightarrow 0} \frac{w^\lambda - 1}{\lambda} &= \lim_{\lambda \rightarrow 0} \frac{e^{\log(w)\lambda} - 1}{\lambda} = \lim_{\lambda \rightarrow 0} \frac{\frac{d}{d\lambda}(e^{\log(w)\lambda} - 1)}{\frac{d}{d\lambda}(\lambda)} = \\ &= \lim_{\lambda \rightarrow 0} \frac{\log(w)e^{\log(w)\lambda}}{1} = \log(w) \end{aligned}$$

We are thus led to consider the natural logarithm function as a utility function:

$$\begin{aligned} U(w) &= \log(w) \\ U'(w) &= \frac{1}{w} = w^{-1} > 0 \\ U''(w) &= -\frac{1}{w^2} = -w^{-2} < 0 \end{aligned}$$

We can rework example 3 using this new utility function:

$$\begin{aligned} f(\alpha) &= 0.5 \log(100 - 0.1\alpha) + 0.5 \log(100 + 0.2\alpha) \\ f'(\alpha) &= \frac{-0.05}{100 - 0.1\alpha} + \frac{0.1}{100 + 0.2\alpha} = 0 \\ \frac{0.1}{100 + 0.2\alpha} &= \frac{0.05}{100 - 0.1\alpha} \\ 0.1(100 - 0.1\alpha) &= 0.05(100 + 0.2\alpha) \\ 0.02\alpha &= 5 \\ \alpha &= 250 \end{aligned}$$

This result agrees exactly with the result we got for $A = 1$ in the equation we derived in example 3.

Thus the natural logarithm function fills the hole quite nicely.

It's interesting to note that, at least in our example, an investor with a logarithmic utility function has very low risk-aversion, since his optimal portfolio is highly leveraged.

7 The Iso-Elastic Utility Functions

All of the utility functions we've examined so far are members of a class called the *iso-elastic utility functions*:

$$U(w) = \begin{cases} \frac{w^\lambda - 1}{\lambda} & \text{for } \lambda < 1, \lambda \neq 0 \\ \log(w) & \text{the limiting case for } \lambda = 0 \end{cases}$$

These functions have the property of *iso-elasticity*, which says that if we scale up wealth by some constant amount k , we get the same utility function (modulo a positive affine transformation). Stated formally,

For all $k > 0$:

$$U(kw) = f(k)U(w) + g(k)$$

for some function $f(k) > 0$ which is independent of w and some function $g(k)$ which is also independent of w .

We can check this formal definition with our utility functions. First consider the case where $\lambda \neq 0$:

$$U(kw) = \frac{(kw)^\lambda - 1}{\lambda} = k^\lambda \left(\frac{w^\lambda - 1}{\lambda} \right) + \frac{k^\lambda - 1}{\lambda} = k^\lambda U(w) + \frac{k^\lambda - 1}{\lambda}$$

Now consider the log function:

$$U(kw) = \log(kw) = \log(k) + \log(w) = U(w) + \log(k)$$

This property of iso-elasticity has a very important consequence for portfolio optimization. It implies that if a given percentage asset allocation is optimal for some current level of wealth, that same percentage asset allocation is also optimal for all other levels of wealth.

We can illustrate this fact by reworking example 3 with initial wealth w_0 a parameter and the amount α invested in the risky asset expressed as a fraction of w_0 . That is, we invest αw_0 dollars in the risky asset and we do nothing with the remaining $(1 - \alpha)w_0$ dollars. In this case the two possible outcomes for ending wealth w are:

$$\begin{aligned} \text{Bad: } w &= 0.9\alpha w_0 + (1 - \alpha)w_0 = (1 - 0.1\alpha)w_0 \\ \text{Good: } w &= 1.2\alpha w_0 + (1 - \alpha)w_0 = (1 + 0.2\alpha)w_0 \end{aligned}$$

The expected utility is:

$$\begin{aligned} f(\alpha) &= 0.5 \frac{(1 - 0.1\alpha)^\lambda w_0^\lambda - 1}{\lambda} + 0.5 \frac{(1 + 0.2\alpha)^\lambda w_0^\lambda - 1}{\lambda} \\ &= \frac{0.5}{\lambda} [(1 - 0.1\alpha)^\lambda w_0^\lambda + (1 + 0.2\alpha)^\lambda w_0^\lambda - 2] \end{aligned}$$

To find the optimal portfolio we take the derivative with respect to α , set it equal to 0, and solve for α :

$$\begin{aligned}
 f'(\alpha) &= \frac{0.5}{\lambda} [-0.1\lambda(1 - 0.1\alpha)^{\lambda-1}w_0^\lambda + 0.2\lambda(1 + 0.2\alpha)^{\lambda-1}w_0^\lambda] = 0 \\
 0.1\lambda(1 - 0.1\alpha)^{\lambda-1}w_0^\lambda &= 0.2\lambda(1 + 0.2\alpha)^{\lambda-1}w_0^\lambda \\
 0.1(1 - 0.1\alpha)^{\lambda-1} &= 0.2(1 + 0.2\alpha)^{\lambda-1} \\
 2 &= \left(\frac{1 - 0.1\alpha}{1 + 0.2\alpha}\right)^{\lambda-1} = \left(\frac{1 + 0.2\alpha}{1 - 0.1\alpha}\right)^{1-\lambda} = \left(\frac{1 + 0.2\alpha}{1 - 0.1\alpha}\right)^A \\
 2^{1/A} &= \frac{1 + 0.2\alpha}{1 - 0.1\alpha} \\
 2^{1/A} - 0.1 \times 2^{1/A}\alpha &= 1 + 0.2\alpha \\
 \alpha &= \frac{2^{1/A} - 1}{0.2 + 0.1 \times 2^{1/A}}
 \end{aligned}$$

This equation for the optimal fraction α to invest in the risky asset is independent of the initial wealth w_0 . Thus the optimal portfolio is the same, regardless of wealth. For example, if the investor's current wealth is \$1,000 and his optimal portfolio is 50% risky and 50% risk-free, then the investor will have exactly the same optimal percentage asset allocation with a current wealth of \$1,000,000.

Thus, investors with iso-elastic utility functions have a constant attitude towards risk expressed as a percentage of their current wealth. This property is called *constant relative risk aversion*.

8 The Negative Exponential Utility Function

Up to this point all of the utility functions we have looked at have been iso-elastic. We now examine a different kind of utility function:

$$\begin{aligned} U(w) &= -e^{-Aw} \quad (\text{for any coefficient of risk aversion } A > 0) \\ U'(w) &= Ae^{-Aw} > 0 \\ U''(w) &= -A^2e^{-Aw} < 0 \end{aligned}$$

This class of utility functions has the interesting property that it is invariant under any translation of wealth. That is:

For any constant k ,

$$U(k+w) = f(k)U(w) + g(k)$$

for some function $f(k) > 0$ which is independent of w and some function $g(k)$ which is also independent of w .

We can easily verify this:

$$U(k+w) = -e^{-A(k+w)} = -e^{-kA}e^{-Aw} = e^{-kA}U(w)$$

To continue our running example, we compute the optimal portfolio in our now familiar investment universe using this new utility function:

$$\begin{aligned} f(\alpha) &= 0.5(-e^{-A(1-0.1\alpha)w_0}) + 0.5(-e^{-A(1+0.2\alpha)w_0}) \\ &= -0.5[e^{-A(1-0.1\alpha)w_0} + e^{-A(1+0.2\alpha)w_0}] \\ f'(\alpha) &= 0.5[0.1Aw_0e^{-A(1-0.1\alpha)w_0} - 0.2Aw_0e^{-A(1+0.2\alpha)w_0}] \\ &= 0.5Aw_0[0.1e^{-A(1-0.1\alpha)w_0} - 0.2e^{-A(1+0.2\alpha)w_0}] \\ &= 0 \\ 0.1e^{-A(1-0.1\alpha)w_0} &= 0.2e^{-A(1+0.2\alpha)w_0} \\ e^{A(1+0.2\alpha)w_0-A(1-0.1\alpha)w_0} &= 2 \\ e^{0.3A\alpha w_0} &= 2 \\ \alpha &= \frac{\log(2)}{0.3} \frac{1}{Aw_0} = \frac{2.31}{Aw_0} \end{aligned}$$

Suppose for sake of example that the coefficient of risk aversion $A = .00231$. Our equation becomes simply $\alpha = 1000/w_0$.

In this example, with a current wealth of \$1,000, the investor's optimal strategy is to invest 100% = \$1,000 in the risky asset and keep back nothing. With a current wealth of \$2,000 his optimal strategy is to invest 50% = \$1,000 in the risky asset and do nothing with the other 50% = \$1,000.

In general, as our investor's wealth increases his portfolio rapidly becomes more conservative. He invests the same absolute amount of money (\$1,000) in the risky asset no matter what his wealth is. For example, with a wealth of

\$1,000,000, he would invest the same \$1,000 in the risky asset and do nothing with \$999,000!

Thus, investors with negative exponential utility functions have a constant attitude towards risk expressed in absolute dollar terms. This property is called *constant absolute risk aversion*.

9 Risk Aversion Functions

We saw earlier that two utility functions are the “same” if they differ by a positive affine transformation:

$$V(w) = aU(w) + b \quad \text{for constants } a > 0 \text{ and } b$$

Differentiate both sides of this equation twice:

$$\begin{aligned} V'(w) &= aU'(w) \\ V''(w) &= aU''(w) \end{aligned}$$

Now divide the second equation by the first equation:

$$\frac{V''(w)}{V'(w)} = \frac{U''(w)}{U'(w)}$$

Conversely, suppose we have any pair of functions U and V for which this last equation holds (the ratio of their second to their first derivatives is the same). Let:

$$f(w) = \frac{V'(w)}{U'(w)}$$

Take the derivative:

$$\begin{aligned} f' &= (V'(U')^{-1})' = V''(U')^{-1} + V'(-1)(U')^{-2}U'' \\ &= \frac{V''}{U'} - \frac{V'U''}{(U')^2} = \frac{V''U' - V'U''}{(U')^2} = \frac{0}{(U')^2} = 0 \end{aligned}$$

$f' = 0$, so we must have $f(w) = a$ for some constant a :

$$\begin{aligned} f(w) &= \frac{V'(w)}{U'(w)} = a \\ V'(w) &= aU'(w) \end{aligned}$$

Integrate both sides of this equation:

$$\begin{aligned} \int V'(w)dw &= \int aU'(w)dw \\ V(w) &= aU(w) + b \quad \text{for some constant } b \end{aligned}$$

We have shown that two utility functions are the “same” if and only if the ratios of their second derivatives to their first derivatives are the same.

For any utility function U this leads us to consider the following function:

$$A(w) = -\frac{U''(w)}{U'(w)}$$

Note that for utility functions we always have $U' > 0$ and $U'' < 0$, so $A(w) > 0$.

This function is called the *Pratt-Arrow absolute risk aversion function*. It completely characterizes the utility function. It provides a measure of the absolute risk aversion of the investor as a function of the investor's wealth.

Let's evaluate this function for the utility functions we've looked at so far.

Iso-elastic, $A =$ coefficient of risk aversion:

$$U(w) = \frac{w^\lambda - 1}{\lambda} : \quad A(w) = -\frac{(\lambda - 1)w^{\lambda-2}}{w^{\lambda-1}} = \frac{1 - \lambda}{w} = \frac{A}{w}$$

$$U(w) = \log(w) : \quad A(w) = -\frac{-w^{-2}}{w^{-1}} = \frac{1}{w} = \frac{A}{w}$$

Negative exponential, $A =$ coefficient of risk aversion:

$$U(w) = -e^{-Aw} : \quad A(w) = -\frac{-A^2e^{-Aw}}{Ae^{-Aw}} = A$$

For the negative exponential functions we have constant absolute risk aversion, as we saw before. For the iso-elastic functions we have decreasing absolute risk aversion.

A related function is the *relative risk aversion function*:

$$A(w)w = -\frac{U''(w)w}{U'(w)}$$

For this function we have:

$$\begin{aligned} \text{Iso-elastic: } A(w)w &= A \\ \text{Negative exponential: } A(w)w &= Aw \end{aligned}$$

Thus for the iso-elastic functions we have constant relative risk aversion. For the negative exponential functions we have increasing relative risk aversion.

10 Example 4 – Variable Risk Aversion

In this example we show how to work backwards from a desired relative risk aversion function to the corresponding utility function. Our target relative risk aversion function is linear and increasing in wealth:

$$A(w)w = -\frac{U''(w)w}{U'(w)} = a + bw \quad \text{for constants } a > 0 \text{ and } b > 0$$

$$wU''(w) + (a + bw)U'(w) = 0$$

Let $V = U'$:

$$wV'(w) + (a + bw)V(w) = 0$$

This first-order ordinary differential equation has the solution:

$$V(w) = w^{-a}e^{-bw}$$

We integrate to get U :

$$U(w) = \int_0^w x^{-a}e^{-bx} dx$$

$$U'(w) = w^{-a}e^{-bw} > 0$$

$$U''(w) = -aw^{-a-1}e^{-bw} - bw^{-a}e^{-bw} = -(a + bw)w^{-a-1}e^{-bw} < 0$$

$$A(w) = -\frac{U''(w)}{U'(w)} = \frac{a}{w} + b$$

$$A(w)w = a + bw$$

Note that we haven't attempted to simplify the integral in the definition of U . It turns out that this isn't even necessary for the purpose of our example.

This utility function has decreasing absolute risk aversion and increasing relative risk aversion. In this sense it is “in between” the iso-elastic family and the negative exponential family.

Once again we use our running example of determining the optimal portfolio with our usual investment opportunity, using the same notation as before.

$$f(\alpha) = 0.5U((1 - 0.1\alpha)w_0) + 0.5U((1 + 0.2\alpha)w_0)$$

$$f'(\alpha) = -0.05w_0U'((1 - 0.1\alpha)w_0) + 0.1w_0U'((1 + 0.2\alpha)w_0) = 0$$

$$U'((1 - 0.1\alpha)w_0) = 2U'((1 + 0.2\alpha)w_0)$$

$$w_0^{-a}(1 - 0.1\alpha)^{-a}e^{-bw_0(1-0.1\alpha)} = 2w_0^{-a}(1 + 0.2\alpha)^{-a}e^{-bw_0(1+0.2\alpha)}$$

$$e^{.3bw_0\alpha} = 2 \left(\frac{1 + 0.2\alpha}{1 - 0.1\alpha} \right)^{-\alpha}$$

$$w_0 = \frac{\log(2) - a \log \left(\frac{1+0.2\alpha}{1-0.1\alpha} \right)}{.3b\alpha}$$



Figure 7: Example 4 – Variable Risk Aversion

As an example we use $a = 2.408$ and $b = 0.000003$. The graph in Figure 7 displays the function above, with α ranging from 0% to 100% displayed on the y axis and initial wealth w_0 displayed using logarithmic scaling on the x axis.

With the values we chose for a and b this investor has low relative risk aversion of 2.438 at wealth \$1,000, increasing linearly to 32.408 at wealth \$10,000,000. Up to around \$10,000 the investor is quite aggressive, placing nearly 100% of his portfolio in the risky asset. After his portfolio has grown in value to \$10,000 he gradually becomes more conservative, with 88% risky at \$100,000, 43% risky at \$1,000,000, and only 8% risky at \$10,000,000.

This investor clearly becomes increasingly concerned about preserving his wealth as his wealth increases. This change is gradual, however, unlike the negative exponential utility functions.

We do not mean to imply that all or even most investors have this characteristic. This is a totally contrived example. Indeed, the same technique can be used to generate any desired pattern of risk aversion as a function of wealth.

11 Utility and Real-Life Investors

What kind of utility function is the best kind to use? What kind of utility function is typical of most investors? What kind best describes the mythical “average” investor?

These, of course, are the \$64,000 questions. There do not appear to be definitive answers. Indeed, it seems that the questions are quite controversial.

We have taken a detailed look at several kinds of utility functions, including the iso-elastic family with constant relative risk aversion, the negative exponential family with constant absolute risk aversion, and an example in between with increasing relative and decreasing absolute risk aversion.

Without some kind of evidence there is no reason to believe that any of these particular utility functions which we have examined describes all investors or even any individual investor or the average investor. It is entirely reasonable for an investor’s attitudes towards risk to vary with the amount of wealth the investor has accumulated, and it’s reasonable for different investors to have different patterns of risk aversion as functions of wealth.

William Sharpe says the following about the notions of constant relative risk aversion and constant absolute risk aversion in reference [2]:

The assumption of constant relative risk aversion seems much closer to the preferences of most investors than does that of constant absolute risk aversion. Nonetheless, it is by no means guaranteed to reflect every Investor’s attitude. Some may wish to take on more risk ... as their wealth increases. Others may wish to take on less. Many analysts counsel a decrease in such risk as one ages. Some strategies are based on acceptance of more or less risk, based on economic conditions. And so on.

For these and other reasons it is important to at least consider strategies in which an Investor’s risk tolerance... changes from time to time. However, such changes, if required at all, will likely be far more gradual than those associated with a constant risk tolerance expressed in terms of end-of-period value.

References

- [1] Robert C. Merton. *Continuous-Time Finance*. Blackwell, 1990.
- [2] William F. Sharpe. Macro-investment analysis (an electronic work-in-progress). <http://www.stanford.edu/~wfsarpe/mia/mia.htm>.