

Utility Theory Under Uncertainty

Lorenzo Naranjo

March 2026

Introduction

In a single period model, agents must decide today how much to consume and how much to save for later. In this notebook we take that decision as given, and assume that agents derive their utility from consumption at the end of the period. As is common in finance, consumption is represented by a single good. Agents prefer more to less, but the marginal utility of each additional unit of consumption is decreasing, i.e., the last bite is never as good as the first one.

We capture this intuition by a utility function $u : \mathbb{R}^+ \rightarrow \mathbb{R}$. The domain of the utility function is real consumption and therefore cannot be negative. In the following, we assume that $u(c)$ is continuous and differentiable of all orders. The level of utility is not important, and can even be negative, but we assume that the utility function is increasing and strictly concave.¹ Mathematically, it must be the case that $u'(c) > 0$ and $u''(c) < 0$. In some cases it will also be useful to consider utility functions such that $u'(0) = \infty$, that is, in starvation an extra unit of consumption provides an infinite amount of extra utility.

A common choice of utility function is

$$u(c) = \frac{c^{1-\gamma} - 1}{1 - \gamma},$$

¹A function $f : D \rightarrow \mathbb{R}$ is concave if

$$f(tx_1 + (1-t)x_2) \geq tf(x_1) + (1-t)f(x_2)$$

for all $x_1, x_2 \in X$ and $0 \leq t \leq 1$. The function is strictly concave if the inequality is strict.

where $\gamma > 0$. A special case occurs when $\gamma = 1$, since by L'Hôpital's rule,²

$$\lim_{\gamma \rightarrow 1} \frac{c^{1-\gamma} - 1}{1 - \gamma} = \lim_{\gamma \rightarrow 1} \frac{-c^{1-\gamma} \ln(c)}{-1} = \ln(c).$$

This function is called **power utility** if $\gamma \neq 1$ and log utility if $\gamma = 1$.³ The marginal utility in this case is given by $u'(c) = c^{-\gamma}$.

Another common utility function is

$$u(c) = -e^{-ac},$$

where $a > 0$, called the **exponential utility** function. We will see that both types of utility functions induce different attitudes under uncertainty.

Expected Utility and Risk Aversion

The analysis in the next two sections closely follows Chapter 1 in Ingersoll (1987). Denote by $\tilde{W}$ the wealth of the agent at the end of the period. This wealth is generated by investing a certain amount today. In general, $\tilde{W}$ is unknown today and can be thought as a random variable defined over a probability space (Ω, P) . We use the tilde on top of W to emphasize that the wealth is unknown today.

When faced with uncertainty, the utility of final consumption is also random. Thus, the agent is not so much worried about the level of consumption she is going to get, but rather she is more worried about the utility she might get from that level of consumption. We can model the agent's behavior using the concept of **expected utility**. The utility derived by random consumption $\tilde{W}$ is given by

$$U(\tilde{W}) = E(u(\tilde{W})).$$

²Both numerator and denominator vanish at $\gamma = 1$, so L'Hôpital's rule applies. Differentiating with respect to γ : $\frac{d}{d\gamma}(c^{1-\gamma} - 1) = -c^{1-\gamma} \ln(c)$ and $\frac{d}{d\gamma}(1 - \gamma) = -1$.

³Many textbooks and programming languages use log instead of ln. In these notebooks I will use ln to denote the natural logarithm.

Thus, we have that $\tilde{W}_1 \succsim \tilde{W}_2$ if $U(\tilde{W}_1) \geq U(\tilde{W}_2)$, and $\tilde{W}_1 \sim \tilde{W}_2$ if $U(\tilde{W}_1) = U(\tilde{W}_2)$. Note that the expected utility of a certain level of wealth W is just $u(W)$.

Consider now a random variable $\tilde{\varepsilon}$ such that $E(\tilde{\varepsilon}) = 0$ and $V(\tilde{\varepsilon}) > 0$. We say that an agent is risk-averse if she prefers a certain level of wealth W over a random payoff with the same expected value. If we define $\tilde{W} = W + \tilde{\varepsilon}$, we have that $E(\tilde{W}) = W$, but $V(\tilde{W}) > 0 = V(W)$. Thus, the agent is risk-averse if

$$u(W) > E(u(W + \tilde{\varepsilon})). \quad (1)$$

In order to understand better the notion of risk-aversion, we will use the following result.

Property 1 (Jensen's Inequality). *Let $f : D \rightarrow \mathbb{R}$ be a twice-continuously differentiable and strictly concave function, and X a random variable defined in a probability space (Ω, P) such that the range of X is contained in the domain of f , and $V(X) > 0$. Then we have that*

$$f(E(X)) > E(f(X)).$$

Proof

Let $m = E(X)$. Since f is strictly concave and differentiable, the tangent line at m lies strictly above the graph of f for all $x \neq m$:

$$f(x) < f(m) + f'(m)(x - m).$$

Setting $x = X$ and taking expectations:

$$E(f(X)) < f(m) + f'(m) E(X - m) = f(m) = f(E(X)),$$

where $E(X - m) = 0$ follows from $m = E(X)$, and the inequality is strict because $V(X) > 0$ ensures $X \neq m$ with positive probability. $\square$

Jensen's inequality shows that strict concavity in u implies risk-aversion. The converse is also true. Consider a risk-averse investor with utility function u and a gamble $\tilde{\varepsilon}$ that pays $(1 - q)a$

with probability q and $-qa$ with probability $1 - q$. The gamble is fair since

$$E(\tilde{\varepsilon}) = q(1 - q)a - (1 - q)qa = 0.$$

A risk-averse investor, though, dislikes the gamble implying

$$u(W) > E(u(W + \tilde{\varepsilon})) = qu(W + (1 - q)a) + (1 - q)u(W - qa). \quad (2)$$

Because $W = q(W + (1 - q)a) + (1 - q)(W - qa)$ for all values of a and q , equation (2) states that u evaluated at a convex combination of two points is larger than the corresponding convex combination of u -values, i.e., u is strictly concave. To see that this specific family of gambles is sufficient for the general case, note that for any x_1, x_2 in the domain of u and any $t \in (0, 1)$, setting $q = t$, $W = tx_1 + (1 - t)x_2$, and $a = x_1 - x_2$ gives $W + (1 - q)a = x_1$ and $W - qa = x_2$, so equation (2) yields $u(tx_1 + (1 - t)x_2) > tu(x_1) + (1 - t)u(x_2)$.

Property 2 (Risk Aversion and Concavity of the Utility Function). *An agent is risk-averse as defined in (1) if and only if her utility function is strictly concave.*

In many situations, we are forced to take on a risky gamble. For example, if you buy a car you face the risk of an accident that can result in costly repairs. Many people choose to pay for insurance and hence reduce this risk. In our framework, we define the insurance premium as the maximum amount that an agent is willing to pay to eliminate the risk.

To formalize this notion, consider an asset with value W . After you buy the asset, you can either face the risk of owning the asset producing a wealth of $W + \tilde{\varepsilon}$, or pay Π_i to insure the asset, which guarantees a certain wealth of $W - \Pi_i$. An agent is indifferent between the two choices if

$$u(W - \Pi_i) = E(u(W + \tilde{\varepsilon})). \quad (3)$$

The value of Π_i that solves (3) is called the **insurance premium**. When comparing two agents' attitudes toward risk, the agent who is willing to pay the most for insurance is **more risk averse** than the other.

In the previous analysis, since the agent is indifferent between the risky gamble and getting $W - \Pi_i$, we call $W - \Pi_i$ the **certainty equivalent** wealth level.

Example 1. Suppose the economy can be in one of the following two states: (i) Boom or “good” state and (ii) Recession or “bad” state, which can occur with equal probability. Consider a risky asset that would have a price of \$50 in the good state and \$10 in the bad state, which currently trades at \$30. Two investors are evaluating this asset.

The utility function of the first investor (A) is

$$U(W) = 10 \ln(W),$$

whereas for the second investor (B) we have

$$U(W) = 2W + 5.$$

What is the maximum price that investors A and B would be willing to pay for the risky asset?

The expected utility of the risky asset for investor A is

$$E(U) = 0.5(10 \ln(10)) + 0.5(10 \ln(50)) = 31.0730.$$

Therefore, the maximum price that the agent is willing to pay for the asset is its certainty equivalent (CE), i.e.,

$$10 \ln(CE) = 31.0730$$

$$\ln(CE) = 3.1073$$

$$CE = \exp(3.1073) = \$22.36.$$

For investor B, the expected utility of the risky asset is

$$E(U) = 0.5(2 \times 10 + 5) + 0.5(2 \times 50 + 5) = 65.$$

The maximum price that the agent is willing to pay for the asset is its certainty equivalent:

$$2 \times CE + 5 = 65$$

$$2 \times CE = 60$$

$$CE = \$30.$$

□

Local Risk Aversion

Intuitively, a function that is more concave should induce more risk aversion than a function that is less concave. We can formalize this intuition by looking at the insurance premium for a gamble with a very small variance.

Let $E(\tilde{\varepsilon}) = 0$ and $V(\tilde{\varepsilon}) > 0$. We know that the insurance premium Π_i solves $u(W - \Pi_i) = E(u(W + \tilde{\varepsilon}))$.

First, do a Taylor expansion of first order of $u(W - \Pi_i)$ around W :

$$\begin{aligned} u(W - \Pi_i) &\approx u(W) + u'(W)(W - \Pi_i - W) \\ &= u(W) - u'(W)\Pi_i. \end{aligned} \tag{4}$$

Second, do a Taylor expansion of second order of $u(W + \tilde{\varepsilon})$ around W :

$$\begin{aligned} u(W + \tilde{\varepsilon}) &\approx u(W) + u'(W)(W + \tilde{\varepsilon} - W) + \frac{1}{2}u''(W)(W + \tilde{\varepsilon} - W)^2 \\ &= u(W) + u'(W)\tilde{\varepsilon} + \frac{1}{2}u''(W)\tilde{\varepsilon}^2 \\ E(u(W + \tilde{\varepsilon})) &\approx u(W) + \frac{1}{2}u''(W)\sigma_{\tilde{\varepsilon}}^2. \end{aligned} \tag{5}$$

Equating (4) and (5) we find that:

$$\Pi_i \approx -\frac{1}{2} \frac{u''(W)}{u'(W)} \sigma_{\tilde{\varepsilon}}^2. \tag{6}$$

The previous expression shows that for an initial wealth of W , the insurance premium depends positively on the local curvature of the utility function at that point as measured by $-u''(W)$.

We denote by

$$ARA = -\frac{u''(W)}{u'(W)}$$

the coefficient of **absolute risk-aversion**, and by

$$RRA = -\frac{u''(W)}{u'(W)}W$$

the coefficient of **relative risk-aversion**.

Intuition for ARA. Equation (6) can be rewritten as

$$\Pi_i \approx \frac{1}{2}ARA \cdot \sigma_\varepsilon^2,$$

so ARA is literally the dollar insurance premium the agent is willing to pay per unit of variance. An agent with a higher ARA will pay more, in absolute dollar terms, to eliminate any given gamble. Because ARA can itself depend on W , it is natural to ask how risk-aversion changes with wealth. An agent with **decreasing absolute risk-aversion** (DARA) requires a smaller dollar premium as wealth grows — risky positions become less burdensome as the agent becomes wealthier. Most empirical evidence and introspection favor DARA.

Intuition for RRA. To see why the factor of W appears, consider a proportional gamble $\tilde{\varepsilon} = W\tilde{\delta}$, where $\tilde{\delta}$ is a percentage shock with $E(\tilde{\delta}) = 0$ and $V(\tilde{\delta}) = \sigma_\delta^2$. Substituting into (6):

$$\Pi_i \approx \frac{1}{2}ARA \cdot W^2\sigma_\delta^2 = \frac{1}{2}RRA \cdot W\sigma_\delta^2.$$

Dividing by W gives the **fraction of wealth** the agent would sacrifice to eliminate the proportional gamble:

$$\frac{\Pi_i}{W} \approx \frac{1}{2}RRA \cdot \sigma_\delta^2.$$

Thus RRA measures the premium per unit variance for proportional gambles, expressed as a share of wealth. An agent with **constant relative risk-aversion** (CRRA) maintains the same fraction of wealth invested in risky assets regardless of her wealth level.

Example 2. Take

$$u(W) = \frac{W^{1-\gamma} - 1}{1-\gamma}$$

Then $u'(W) = W^{-\gamma}$ and $u''(W) = -\gamma W^{-\gamma-1}$, implying that

$$RRA = -\left(\frac{-\gamma W^{-\gamma-1}}{W^{-\gamma}}\right)W = \gamma$$

Power utility is an example of a function that exhibits constant relative risk-aversion. $\square$

Example 3. For exponential utility $u(W) = -e^{-aW}$ with $a > 0$, we have $u'(W) = ae^{-aW}$ and $u''(W) = -a^2e^{-aW}$, so

$$ARA = \frac{a^2e^{-aW}}{ae^{-aW}} = a.$$

The coefficient of absolute risk-aversion is constant (equal to a) and independent of W . Exponential utility therefore exhibits **constant absolute risk-aversion** (CARA). A direct implication is that the insurance premium $\Pi_i \approx \frac{1}{2}a\sigma_\varepsilon^2$ does not depend on the agent's wealth: a billionaire and a middle-class worker with the same CARA parameter a would pay identical dollar amounts to insure against the same gamble. This is analytically convenient but empirically implausible, since most people become less risk-averse in absolute terms as they accumulate wealth. $\square$

References

Ingersoll, Jonathan E. 1987. *Theory of Financial Decision Making*. Vol. 3. Rowman & Littlefield.