23. Attitudes Toward Risk

Attitudes toward risk are an important topic of study concerning choice under uncertainty. Sometimes people are risk averse, buying insurance, refusing to bet in casinos or play the lottery, and generally avoiding risk. Others embrace risk, betting on long odds, doubling down, or maybe simply taking a small risk on the lottery. Some do both at the same time.

These attitudes toward risk can change depending on the economic situation. Expected utility provides a framework where attitudes toward risk can be quantified, where we can study their effect on economic decision-making more precisely, and where such attitudes can change with the economic situation.

Section one investigates several basic concepts of risk aversion—concavity, certainty equivalence, risk premium, probability premium—and shows how they are related. We take a short break to consider an important issue. Can it make sense for the same person to both gamble and insure? Section three applies these risk concepts to a variety of economic problems. Section four quantifies attitudes toward risk via the Arrow-Pratt measures of risk aversion. We find that changes in the amount of risk aversion can affect economic behavior. The first four sections all consider the effect of risk on economic agents. Section five looks at risk itself, at the question of when one random variable is riskier than another. We go beyond consideration of means and variances to consider stochastic dominance and mean-preserving spreads. In section six, we learn a bit about the limitations of the expected utility approach. Certain kinds of behavior that seems to make sense has been ruled out. The examples there are the beginnings of alternative theories of consumer and producer choice under uncertainty.

Outline:

1. Risk Aversion
2. Why Do People Both Gamble and Insure?
3. Applications of Expected Utility
4. Arrow-Pratt Risk Aversion
5. Stochastic Dominance
6. The Limits of Expected Utility

23.1 Risk Aversion

The basic idea of risk aversion is that risk averse agents prefer to avoid risk while risk lovers seek risk. A risk averse agent will prefer to buy fair insurance rather than bearing the risk himself. A risk loving agent prefers to gamble, even if the odds slightly favor the house. We can formalize these notions by comparing asking whether the agent prefers a lottery or a certain outcome with the same expected value.

We say an agent who prefers to take the expected value of a lottery rather than the lottery itself is *risk averse*. That is, if $EX \succsim X$. An agent is *risk neutral* if $EX \sim X$ and *risk seeking* or *risk loving* if $X \succsim EX$.

23.1.1 Quadratic Utility

A simple method of modeling risk averse and risk loving behavior is to use preferences that depend only on the mean μ and variance σ^2 of a random variable X.[1] Risk aversion then translates into variance being a bad, while risk lovers consider variance a good. Such preferences can be rationalized by use of a quadratic expected utility function.

Example 23.1.1: Quadratic Utility. Suppose $u(x) = a + bx - cx^2$ with $a, b, c > 0$. This utility function is risk averse due to the negative sign on x^2. Let X be a random variable with mean μ and variance σ^2. Expected utility of a lottery X is

$$\begin{aligned} Eu(X) &= a + bEX - cE(X^2) \\ &= a + b\mu - c\mu^2 - c\sigma^2 \end{aligned}$$

because $E(X^2) = \sigma^2 + \mu^2$.

Using the formula for $Eu(X)$, we can rewrite expected utility as

$$Eu(X) = u(\mu) - c\sigma^2 = u(EX) - c\sigma^2.$$

This makes it clear that for any lottery that is not certain ($\sigma^2 > 0$), $Eu(X) < u(EX)$. In other words, the consumer is risk averse.

The case $c = 0$ represents risk neutrality, while $c < 0$ means the consumer is risk loving.

With this utility function, the effect of an change in mean is ambiguous. However, if we restrict our attention to distributions supported on the increasing portion of u ($x < b/2c$), an increase in μ increases expected utility. ◀

[1] E.g., Markowitz's (1952a) well-known model of portfolio selection.

23.1.2 Risk Aversion and Concavity

A key result is that concave utility indicates risk aversion while convex utility translates to risk seeking as in Figure 23.1.3).[2]

Theorem 23.1.2. *A utility function $u\colon A \to \mathbb{R}$ where $A \subset \mathbb{R}$ is risk averse if u is concave, risk neutral if u is affine, and risk seeking if u is convex.*

Proof. By Jensen's inequality

$$\int u(x)\, dF(x) \left\{ \begin{matrix} \leq \\ = \\ \geq \end{matrix} \right\} u\left(\int x\, dF(x)\right) \text{ if } u \text{ is } \left\{ \begin{matrix} \text{concave} \\ \text{affine} \\ \text{convex} \end{matrix} \right\} \text{ in } x.$$

The definitions now yield the result. □

[2] Friedman and Savage (1948) noted this using diagrams similar to Figure 23.1.5.

23.1.3 Risk Aversion in Graphs

We can use Theorem 23.1.2 to illustrate risk aversion by considering the graph of utility. The expected utility of a lottery over two payoffs a and b is then on the line segment between the corresponding points on the utility curve. It corresponds to the point directly over the expected value of the lottery. Figure 23.1.3a illustrates the risk averse case, when concavity ensures that $u(EL)$ lies above $Eu(L)$. This says that the consumer prefers a riskless payoff of EL to taking the risky lottery.

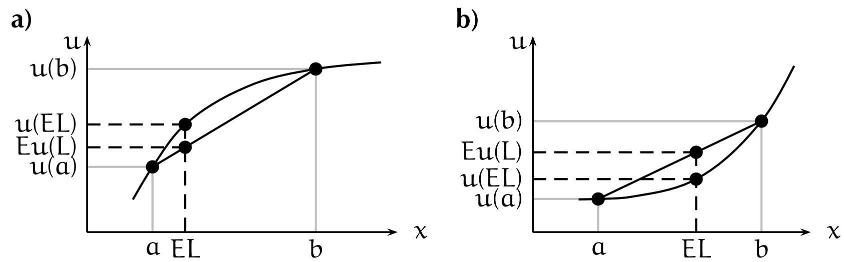

Figure 23.1.3: Risk Aversion and Concavity/Convexity. Let $L = (1-p)a \oplus pb$, so $EL = (1-p)a + pb$.

In panel (a), utility is concave. Then the point $\big((1-p)a + pb, (1-p)u(a) + pu(b)\big) = \big(EL, Eu(L)\big)$ lies below the point $\big(EL, u(EL)\big)$. The fact that $\big(EL, u(EL)\big)$ lies above $\big(EL, Eu(L)\big)$ implies the consumer is risk averse.

Utility is convex in panel (b). In that case, the point $\big((1-p)a + pb, (1-p)u(a) + pu(b)\big) = \big(EL, Eu(L)\big)$ lies above the point $\big(EL, u(EL)\big)$. Then $Eu(L)$ is higher than $u(EL)$, indicating that the consumer is risk seeking.

To see how this works, consider the line segment from $\big(a, u(a)\big)$ to $\big(b, u(b)\big)$. Because u is concave, this lies below the graph of u. Points on that segment have the form

$$\big((1-p)a + pb, (1-p)u(a) + pu(b)\big)$$

for $0 \le p \le 1$. Now consider the lottery L paying a with probability $(1-p)$ and b with probability p. Then

$$\big((1-p)a + pb, (1-p)u(a) + pu(b)\big) = \big(EL, Eu(L)\big).$$

The fact that $\big(EL, u(EL)\big)$ lies above $\big(EL, Eu(L)\big)$ tells us the consumer prefers the safe payoff EL to the risky L—the consumer is risk averse.

Figure 23.1.3b shows the case where u is convex. In that case, $u(EL)$ lies below $Eu(L)$. Such a consumer is a risk seeker who prefers to take a risk rather than settle for a sure thing.

23.1.4 Mean-Preserving Spread

Another way to illustrate risk aversion uses a Marschak-Machina triangle where we add risk by using a mean-preserving spread.

Mean-preserving Spread. A lottery L' is a *mean-preserving spread* of L if the means are the same ($EL' = EL$) and $\text{var}(L') > \text{var}(L)$.

23.1.5 Mean-Preserving Spread and Risk Aversion

For the Marschak-Machina triangle, we take a mean-preserving spread about s_2. This allows us to increase risk without changing the expected value of the lottery. This way we can focus on the attitude toward risk alone without contaminating our comparisons with changes in expected value.

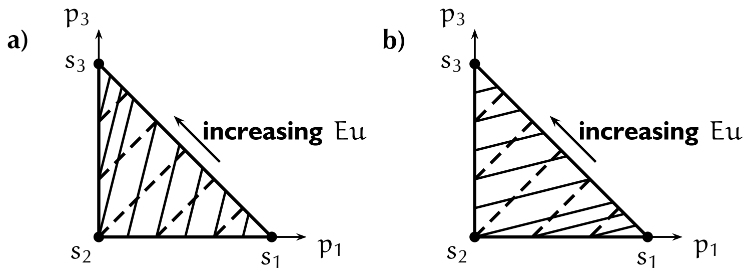

Figure 23.1.4: Risk Preference in the Marschak-Machina triangle. In both panels, the upward sloping solid lines are indifference curves for a case with $s_1 \prec s_2 \prec s_3$. The dashed lines are points of equal expected value, but with increasing risk as we move outward.

In panel (a), you'll notice that lotteries further right on the iso-expected value line *lower* the consumer's utility. In other words, this consumer is risk averse.

Panel (b) is the opposite. Lotteries further right on the iso-expected value line *raise* the consumer's utility. In other words, this consumer is risk seeking.

Let v_s be the value of state s, and suppose $v_1 = v_2 - b$ and $v_3 = v_2 + b$ where $v_2 > b > 0$. This way, the order of the values matches the order of preference: $v_3 > v_2 > v_1$. Now consider the 45° lines in the Marschak-Machina triangle. The probabilities obey $p_3 = a + p_1$ for some $a \in [-1, +1]$. Consider the lottery

$$\begin{aligned} L &= p_1 v_1 \oplus p_2 v_2 \oplus p_3 v_3 \\ &= p_1 v_1 \oplus \big(1 - p_1 - (a + p_1)\big) v_2 \oplus (a + p_1) v_3. \end{aligned}$$

This has expected value

$$p_1(v_2 - b) + (1 - a - 2p_1)v_2 + (a + p_1)(v_2 + b) = v_2 + ab.$$

It is constant on the 45° lines and positive since $v_2 + ab \geq v_2 - b > 0$.

This means that a movement up and right along any 45° line leaves the expected value unchanged, but increases the variance, adding an increasing amount of risk. If utility decreases as the consumer moves out along the iso-expected value lines, the consumer is risk averse. If utility increases, the consumer is risk seeking.

23.1.6 Certainty Equivalent

There are several other ways to measure attitude toward risk. One particularly useful method is the certainty equivalent.

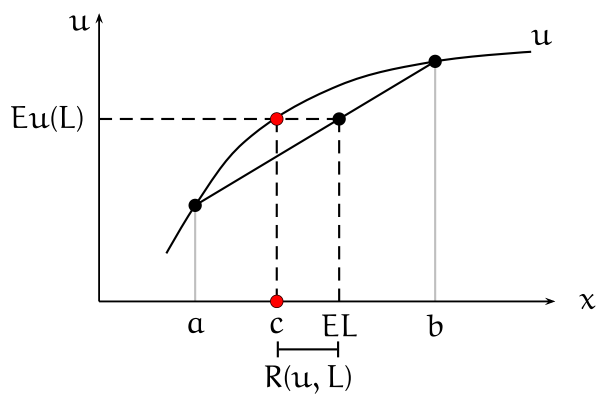

Figure 23.1.5: Certainty Equivalent and Risk Premium. We consider the same setting as in Figure 23.1.3. The lottery with expected value EL results in expected utility $Eu(L)$. We find the intersection with the utility curve u and drop down to determine the certainty equivalent c (see the red dots). That is, we find the c with $u(c) = Eu(L)$.

The consumer receives the same utility from c with certainty as he or she does from the lottery L over a and b with expected payoff EL. Concavity guarantees that the certainty equivalent c lies to the left of EL. It is easily verified that when utility is convex, the certainty equivalent lies to the right of EL.

We can also find the risk premium $R(u, L)$ on this graph by taking the difference between c and EL.

Certainty Equivalent. We say $c(u, L)$ is the certainty equivalent of a lottery L if a person with utility function u is indifferent between the lottery L and receiving $c(u, L)$ with certainty. In other words, if $u\big(c(u, L)\big) = Eu$.

When u is unambiguous, we will sometimes use the shorthand notation c_L to denote the certainty equivalent $c(u, L)$.

23.1.7 Risk Premium

We use the certainty equivalent to define the *risk premium* $R(u, L)$. It is the difference between the expected value of L and the certainty equivalent,[3]

$$R(u, L) = EL - c(u, L).$$

It is also illustrated in Figure 23.1.5.

By Theorem 23.1.10 below, the risk premium is positive if the consumer is risk averse, zero if risk neutral, and negative if risk seeking. If positive, the risk premium tells us how much certain consumption the lottery costs, compared to its expected value.

[3] The term "risk premium" is used in several different ways in the literature.

23.1.8 Example: Certainty Equivalent

Let's compute the certainty equivalent and risk premium.

Example 23.1.6: Certainty Equivalent. Suppose utility is $u(x) = x^{1/3}$ and a lottery X has distribution function $F(x) = 0$ for $x < 0$ or $x > 1$ and $F(x) = x$ on $[0, 1]$. The expected value is $EX = 1/2$ and expected utility is

$$Eu(X) = \int_0^1 x^{1/3}\,dx = \frac{3}{4}x^{4/3}\Big|_0^1 = \frac{3}{4}.$$

This consumer is risk averse because $3/4 < u(EX) = (1/2)^{1/3}$ (this also follows from concavity of u).

The certainty equivalent $c = c(X)$ is found by setting $u(c) = 3/4$ or $c^{1/3} = 3/4$. Then $c(X) = 27/64$ is the certainty equivalent.

It follows that $R = EX - c(X) = 1/2 - 27/64 = 5/64$ is the risk premium. If the consumer owns this lottery, they are willing to sell it for a price as low as $5/64$ units of consumption below its expected value. ◀

23.1.9 The Probability Premium 03/31/22

We consider one final measure of risk aversion: the probability premium. This is defined relative to income x. It asks how much the probability would have to deviate from 50-50 to make a spread of $x - \varepsilon$ and $x + \varepsilon$ indifferent to x itself.[4] The idea is that if you are risk averse more probability must be put on the higher outcome, while if you are risk seeking more probability must be put on the lower outcome.

Probability Premium. For $\varepsilon > 0$, the *probability premium* $\pi(x, \varepsilon, u)$ is defined as the solution π to

$$u(x) = \left(\tfrac{1}{2} + \pi\right) u(x + \varepsilon) + \left(\tfrac{1}{2} - \pi\right) u(x - \varepsilon). \tag{23.1.1}$$

For $\varepsilon = 0$, define $\pi(x, 0, u) = 0$.

We can rearrange equation 23.1.1 to derive an expression for the probability premium.

$$\pi(x, \varepsilon, u) = \frac{u(x) - \frac{1}{2}u(x + \varepsilon) - \frac{1}{2}u(x - \varepsilon)}{u(x + \varepsilon) - u(x - \varepsilon)}. \tag{23.1.2}$$

Although $\pi(x, 0, u)$ is undefined, we can define it as the limit of the right-hand side of equation 23.1.2 as $\varepsilon \to 0$. Plugging in $\varepsilon = 0$ fails as we obtain $0/0$. Because of this, we compute the limit using l'Hôpital's rule. Provided $u'(x) \neq 0$, that yields $\pi(x, 0, u) = 0$.

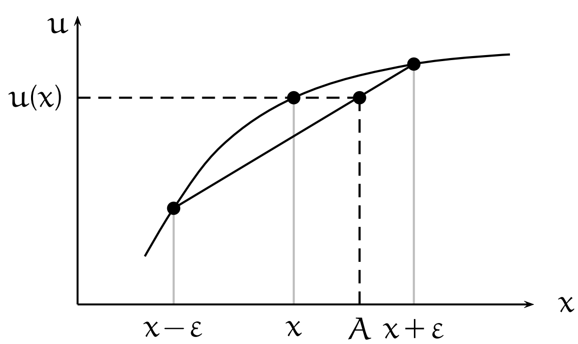

Figure 23.1.7: Probability Premium. The point $A = x + 2\pi\varepsilon$ is the expected payoff when the probability premium is π. It is found by taking the line segment between $\big(x - \varepsilon, u(x - \varepsilon)\big)$ and $\big(x + \varepsilon, u(x + \varepsilon)\big)$, and finding the point where the height is $u(x)$. The corresponding point A on the x-axis gives us the probability premium by inverting the above formula: $\pi = (A - 2)/2\varepsilon$. Concavity ensures $A > x$. Convexity would imply $A < x$.

[4] This is a mean-altering spread when the probabilities are not 50-50.

23.1.10 Example: Probability Premium

Example 23.1.8: Probability Premium. Let $u(x) = -e^{-x}$. Notice that $u' > 0$ and $u'' < 0$. Set

$$-e^{-x} = -\left(\frac{1}{2} + \pi\right) e^{-x-\varepsilon} - \left(\frac{1}{2} - \pi\right) e^{-x+\varepsilon}$$

Factoring out $-e^{-x}$, we obtain

$$1 = \left(\frac{1}{2} + \pi\right) e^{-\varepsilon} + \left(\frac{1}{2} - \pi\right) e^{+\varepsilon}.$$

In this case the probability premium is independent of x. It is

$$\begin{aligned}
\pi(x, \varepsilon, u) &= \frac{2 - e^{-\varepsilon} - e^{+\varepsilon}}{2(e^{-\varepsilon} - e^{+\varepsilon})} \\
&= -\frac{(e^{+\varepsilon/2} - e^{-\varepsilon/2})^2}{2(e^{-\varepsilon/2} - e^{+\varepsilon/2})(e^{-\varepsilon/2} + e^{+\varepsilon/2})} \\
&= \frac{1}{2}\left(\frac{e^{+\varepsilon/2} - e^{-\varepsilon/2}}{e^{+\varepsilon/2} + e^{-\varepsilon/2}}\right) \\
&= \frac{1}{2}\tanh\left(\frac{\varepsilon}{2}\right)
\end{aligned}$$

where tanh is the hyperbolic tangent. ◀

23.1.11 The Derivative of the Probability Premium I

The derivative of the probability premium at $\varepsilon = 0$ can be expressed in terms of the derivatives of the utility function. The expression on the right hand side of equation 23.1.3 is called the *absolute risk aversion*.

Lemma 23.1.9. *Suppose* $u \in \mathcal{C}^2$ *with* $u' > 0$. *Then*

$$4\frac{\partial \pi}{\partial \varepsilon}(x, 0, u) = -\frac{u''(x)}{u'(x)}. \tag{23.1.3}$$

Proof. We start by repeating equation 23.1.2,

$$\pi(x, \varepsilon, u) = \frac{u(x) - \frac{1}{2}u(x+\varepsilon) - \frac{1}{2}u(x-\varepsilon)}{u(x+\varepsilon) - u(x-\varepsilon)}. \tag{23.1.2}$$

Since u is increasing, the denominator is positive for $\varepsilon > 0$. It follows that π is $\mathcal{C}^2$ in (x, ε) for $\varepsilon > 0$.

Proof continues. . .

23.1.12 The Derivative of the Probability Premium II

Rest of Proof. To calculate the derivative at zero, it is easier to consider

$$\pi(x, \varepsilon, u)\big[u(x+\varepsilon) - u(x-\varepsilon)\big] = u(x) - \frac{1}{2}u(x+\varepsilon) - \frac{1}{2}u(x-\varepsilon).$$

The first derivative yields

$$\begin{aligned}\frac{\partial \pi}{\partial \varepsilon}\big[u(x+\varepsilon) - u(x-\varepsilon)\big] &+ \pi\big[u'(x+\varepsilon) + u'(x-\varepsilon)\big] \\ &= -\frac{1}{2}[u'(x+\varepsilon) - u'(x-\varepsilon)].\end{aligned}$$

Letting $\varepsilon \to 0$, would kill off the $\partial\pi/\partial\varepsilon$ term, so we differentiate a second time, obtaining

$$\begin{aligned}\frac{\partial^2 \pi}{\partial \varepsilon^2}\big[u(x+\varepsilon) - u(x-\varepsilon)\big] &+ 2\frac{\partial \pi}{\partial \varepsilon}\big[u'(x+\varepsilon) + u'(x-\varepsilon)\big] \\ &+ \pi\big[u''(x+\varepsilon) - u''(x-\varepsilon)\big] \\ &= -\frac{1}{2}[u''(x+\varepsilon) + u''(x-\varepsilon)].\end{aligned}$$

Finally, let $\varepsilon \to 0$, to obtain

$$4\frac{\partial \pi}{\partial \varepsilon}(x, 0, u)\, u'(x) = -u''(x).$$

Divide by $u'(x)$ to get equation 23.1.3. □

23.1.13 Equivalence of Risk Aversion Concepts

It turns out that all of these notions of attitudes toward risk are the same!

Theorem 23.1.10. *The following are equivalent for a $\mathcal{C}^2$ expected utility function u with $u' > 0$:*

1. u *is risk averse (risk seeking).*
2. u *is concave (convex).*
3. $c(u, F) \le EF$ *for all* F *(*$c(u, F) \ge EF$ *for all* F*).*
4. $R(u, F) \ge 0$ *for all* F *(*$R(u, F) \le 0$ *for all* F*).*
5. $\pi(x, \varepsilon, u) \ge 0$ *(*$\pi(x, \varepsilon, u) \le 0$*) for all* x, ε.

Proof. Jensen's inequality shows (1) and (2) are equivalent, while (3) is just a rewriting of (1). The definition of risk premium shows that (3) and (4) are equivalent. Thus (1)-(4) are equivalent.

As for (5), if $\pi = \pi(x, \varepsilon, u)$, the consumer is indifferent between x and a lottery L with expected value $x + 2\pi\varepsilon$. If the consumer is risk averse, $Eu(L) \le u(EL) = u(x + 2\pi\varepsilon)$. But $u(x) = Eu(L)$, so $u(x) \le u(x + 2\pi\varepsilon)$. This implies $\pi \ge 0$, establishing (5) for risk averse consumers. The risk loving case is similar. It follows that any of (1)-(4) imply (5).

The hard part is showing (5) implies one of the others. Fortunately, the hard work was already done in Lemma 23.1.9. Now for $\varepsilon > 0$,

$$\frac{\pi(x, \varepsilon, u)}{\varepsilon} = \frac{\pi(x, \varepsilon, u) - \pi(x, 0, u)}{\varepsilon} \ge 0.$$

Let $\varepsilon \to 0$ and use l'Hôpital's rule to obtain $-u''/4u' \ge 0$. It follows that $u'' \le 0$, so u is concave. $\square$

23.1.14 Measuring Risk: Risk Premium

One important way to measure risk is via the risk premium, the difference between the expected value of a lottery and its certainty equivalent. The risk premium can be approximated by using the absolute risk aversion, $r_A(x) = u''(x)/u'(x)$.

Example 23.1.11: Risk Premium. Recall that the risk premium is $R(u, F) = \mathrm{E}F - c(u, F)$, so $c(u, F) = \mathrm{E}F - R(u, F)$. This means that

$$u\big(c(u, F)\big) = u(\mathrm{E}F - R) = \mathrm{E}u(F). \tag{23.1.4}$$

When R is small, we may approximate $u(\mathrm{E}F - R)$ by $u(\mathrm{E}F) - u'(\mathrm{E}F)R$. For the right-hand side, we use the second-order Taylor approximation (partly because x may vary significantly from $\mathrm{E}F$). Thus

$$\begin{aligned} \mathrm{E}u(F) &\approx \int \left[u(\mathrm{E}F) + u'(\mathrm{E}F)(x - \mathrm{E}F) + \frac{1}{2}u''(\mathrm{E}F)(x - \mathrm{E}F)^2\right] dF(x) \\ &= u(\mathrm{E}F) + 0 + \frac{1}{2}u''(\mathrm{E}F)\,\mathrm{var}(F). \end{aligned}$$

Substituting the approximations in equation 23.1.4 and eliminating the common term $u(\mathrm{E}F)$, we find $-u'(\mathrm{E}F)R \approx u''(\mathrm{E}F)\,\mathrm{var}(F)/2$. Then the risk premium is

$$R \approx -\frac{u''(\mathrm{E}F)}{u'(\mathrm{E}F)}\frac{\mathrm{var}(F)}{2}.$$

We previously encountered the expression $-u''/u'$ in Lemma 23.1.9. It is called the *coefficient of absolute risk aversion* and denoted $r_A(x) = -u''/u'$. We can use it to write the risk premium as

$$R \approx r_A(\mathrm{E}F)\,\frac{\mathrm{var}(F)}{2}.$$

The price of risk is the ratio of the risk premium to the expected payoff

$$\frac{R}{\mathrm{E}F} \approx r_A(\mathrm{E}F)\,\frac{\mathrm{var}(F)}{2\mathrm{E}F}.$$

Not surprisingly, a higher variance leads to a bigger risk premium for risk averse consumers. ◀

23.1.15 Risk Spreading

Two important applications of the risk premium are risk pooling and risk spreading. Risk spreading refers to cutting up a single, usually large, risky payoff and spreading the risk across many individuals. Risk pooling is the opposite, where many risks are pooled together, with each person taking a share of the aggregate payoff.

As shown in Example 23.1.12, this reduces the aggregate cost of risk. One implication is that firms whose stock is widely held will face a low cost of risk and act much like a risk neutral firm.

Example 23.1.12: Risk Spreading. The cost of risk can be reduced by spreading it among many people. Suppose we divide a risky payoff with distribution F equally among N people. Each will then receive a payoff with mean EF/N and variance $\mathrm{var}(F)/N^2$. If each has wealth w, the risk premium required by each to take the risk will be about $r_A(w)\,\mathrm{var}(F)/2N^2$. The total cost of the risk will be N times that, about

$$r_A(w)\,\frac{\mathrm{var}(F)}{2N}.$$

Thus the total cost of the risk goes to zero as the number of people it is spread among goes to infinity. ◀

23.1.16 Risk Pooling

Risk pooling involves diversification of risk and is the basis for mutual insurance companies. A mutual insurance company is owned entirely by its policyholders. In essence, they are insuring each other. Risk pooling makes this possible. In risk pooling, a group of people faces many lotteries (although perhaps only a few each). The lotteries are pooled to form an aggregate payoff. Each person takes share of this pooled lottery. Example 23.1.13 shows that risk pooling can yield a payoff that is much less risky than any of the individual payoffs, with a corresponding reduction in the cost of risk.

Example 23.1.13: Risk Pooling. Suppose we have N people, with person n receiving a risky lottery x_n. The lotteries are iid with distribution function F. The N people decide to pool their risk. They add together the payoffs and each take an equal share. Each person's payoff is $(1/N)\sum_n x_n$, which has mean EF and variance var(F)/N. It follows that the risk premium for person n is about

$$r_A(\mathrm{EF})\,\frac{\mathrm{var(F)}}{2N}.$$

This is (1/N) times the risk premium they would face if they did not pool. ◀

Risk pooling is also the basis of annuities, which can provide a way of insuring against outliving your assets. In that case the pooling works a little differently. The assets of many people are put in a pool. The assets earn returns (and losses). Income is paid to the survivors each year, based on the life expectancy of the remaining pool.

23.2 Why Do People Both Gamble and Insure?

We take a moment to return to the basic concepts and asking ourselves why some people gamble, and others insure. And what about those who do both at the same time? How does this relate to risk aversion. If one group always liked to take fair gambles, and another always preferred to buy fair insurance, we could simply say that the first group consists of risk lovers and the second consists of risk averters. But the people who do both would be a problem.

23.2.1 Friedman and Savage's Answer

Friedman and Savage (1948) proposed a solution to the problem that the same people buy both lottery tickets and insurance. They suggested the solution lay in the form of the utility function.

Example 23.2.1: Lottery Tickets and Insurance. Suppose utility is concave at low income levels and convex at higher income levels. Such a utility function is shown in Figure 23.2.2. It has an inflection point at I^*, and is risk averse at any smaller level of consumption. Any consumer has income below I^* will be happy to purchase insurance against large losses even if the price is a bit above the actuarially fair level.

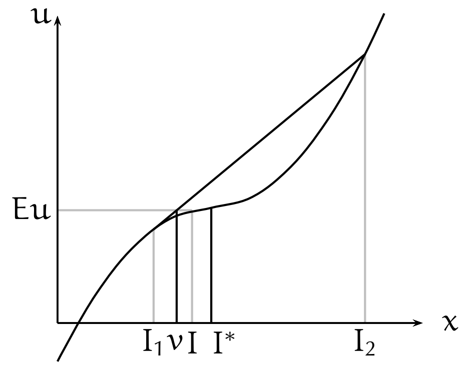

Figure 23.2.2: Friedman-Savage Utility. Utility has in inflection point at I^*. If income is less than I^*, the consumer is risk averse as long as the income involved is never too large. Such a consumer would be willing to insure against a large loss. However, such a consumer will also seek risks that promise a large gain. Thus if income is I, a gamble that offers a possible loss to I_1 and gain to I^2 will be taken if the expected loss is small (less than $I - v$.

A consumer with income I would also be willing to play the lottery if the gains were large enough. Here a lottery ticket costing $I - I_1$ and potentially paying I_2 is preferred to remaining with certain income I provided the expected value of the lottery is at least v.

This is not the end of the story. Friedman and Savage (1948) also argued that utility would have to be concave at high income levels (not shown on Figure 23.2.2). The utility curve must have a second inflection point. Bounded utility is used to avoid the St. Petersburg paradox. ◄

23.2.2 Markowitz's Answer

Markowitz (1952b) considered the Friedman-Savage example further. He examined how consumers would react to various lotteries if they had such utility. He found that middle income individuals would take large symmetric fair bets.[5] Yet this is rarely observed. He also found that near-rich individuals would not only not insure, but would be willing to underwrite insurance, even at substantial cost. This too is rarely observed, if ever.

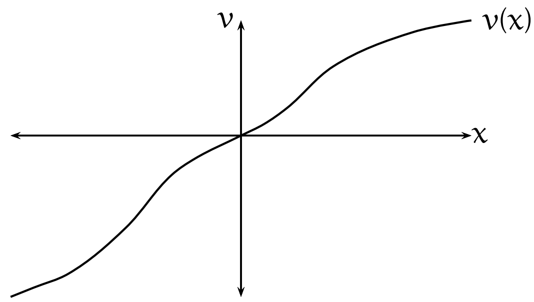

Figure 23.2.3: Markowitz Utility.

Markowitz then suggested writing utility relative to current wealth. In other words, there is a function v so that utility is $u(x) = v(x-x_0)$ with x_0 denoting current wealth. This should be bounded below and above (no St. Petersburg paradox) with three inflection points, one at current wealth, one above, and one below. He also argued that the v should average steeper to the left of the origin that to the right because people generally avoid symmetric bets. This is illustrated in Figure 23.2.3.[6]

[5] You may find Markowitz's (1952b) discussion based on his Figure 3 to be confusing. This is because he refers to points C and D that are not present on the graph. He really means A and B.
[6] The similarity to portions of Kahneman and Tversky's (1979) prospect theory is striking.

23.3 Applications of Expected Utility

03/25/21

We consider several applications of expected utility in this section. Some applications inherently involve uncertainty, as with insurance and demand for risky assets.

23.3.1 Uncertain Price and the Risk Neutral Firm

One easy application of expected utility concerns price-taking firms facing an uncertain price. We start by considering the risk neutral firm, then follow with the risk averse firm.

Example 23.3.1: Uncertain Price and the Risk Neutral Firm. Suppose a firm is a price-taker, but must decide how much to produce before the market price p is known. The firm knows the price distribution function $F(p)$ and cost as a function of output, $C(q)$. We presume C is twice continuously differentiable with $C' > 0$ and $C'' > 0$. In this case, the firm is risk neutral, so it maximizes expected profits

$$E\pi(q) = \int \big[pq - C(q)\big]\, dF(p) = q \int p\, dF(p) - C(q).$$

The first-order condition is

$$\int p\, dF(p) = C'(q)$$

and it is easily verified that the second-order sufficient condition holds, that $-C''(q) < 0$. The first-order condition can be rewritten $Ep = C'(q)$. Rather than setting marginal cost equal to price, the risk neutral firm sets marginal cost equal to **expected** price. ◀

23.3.2 Uncertain Price and the Risk Averse Firm I

To study the risk averse firm, we use the same framework as in Example 23.3.1. The only change is that we assume the firm to be risk averse and maximize the expected utility of profit. Using expected utility was not necessary in the risk neutral case as maximizing expected utility and expected profit are the same under risk neutrality.

Example 23.3.2: Uncertain Price and the Risk Averse Firm. The risk averse firm maximizes expected utility of profit, which is

$$Eu(F) = \int u\big(pq - C(q)\big)\,dF(p).$$

The first-order condition is

$$\begin{aligned} 0 &= \int u'\big(pq - C(q)\big)\big[p - C'(q)\big]\,dF(p) \\ &= \int u'(\pi)\big[p - C'(q)\big]\,dF(p), \end{aligned}$$

which can be written

$$\begin{aligned} E\big[pu'(\pi)\big] &= E\big[C'(q)u'(\pi)\big] \\ &= C'(q)\,E\big[u'(\pi)\big]. \end{aligned}$$

It is easy to see that this reduces to the risk neutral case when $u(x) = ax + b$.

23.3.3 Uncertain Price and the Risk Averse Firm II

We claim that

$$\left[u'(\pi) - u'\big(E\pi\big)\right]\left[p - Ep\right] \leq 0.$$

To establish the claim, we employ the argument of Sandmo (1971). First suppose $p > Ep$. Then

$$\pi = pq - C(q) > (Ep)q - C(q) = E\pi$$

so $u'(\pi) < u'(E\pi)$ because $u'' < 0$. Similarly, when $p < Ep$, $u'(\pi) > u'(E\pi)$. It follows that

$$\left[u'(\pi) - u'(E\pi)\right]\left[p - Ep\right] \leq 0, \tag{23.3.5}$$

This is the desired inequality and establishes the claim.

Now rewrite equation 23.3.5 as

$$u'(\pi)(p - Ep) \leq u'(E\pi)(p - Ep)$$

and take expectations to obtain

$$E\left[u'(\pi)\big(p - Ep\big)\right] \leq E\left[u'(E\pi)\big(p - Ep\big)\right] = E\left[u'(\pi)\right]E\left[p - Ep\right] = 0,$$

so $E[pu'(\pi)] \leq E[p]\,E[u'(\pi)]$.

It follows that

$$\begin{aligned} C'(q)\,E[u'(\pi)] &= E[pu'(\pi)] \\ &\leq E[p]\,E[u'(\pi)] \end{aligned}$$

where the first line is the first-order condition and the second follows by Sandmo's argument.

Since $u' > 0$, $E[u'] > 0$, implying $C'(q) \leq Ep$. Let $\bar{q}$ be the quantity the risk neutral firm would produce. From Example 23.3.1, $Ep = C'(\bar{q})$, so $C'(q) \leq C'(\bar{q})$. Since $C'' > 0$, the optimal $q \leq \bar{q}$. The presence of uncertainty reduces output.

Sandmo also shows that output decreases when fixed cost increases if and only if the absolute risk aversion is decreasing. ◄

23.3.4 Savings under Uncertainty SKIPPED

We can also consider consumer's problems where prices are uncertain. One natural home for this is a consumption-savings model where the interest rate is uncertain.

Example 23.3.3: Savings under Uncertainty. We use a two-period model to investigate savings under uncertainty. Let c_t denote consumption in period $t = 1, 2$. In period 1, the consumer receives wealth w which can be divided between consumption in period 1, c_1, and savings, $(w - c_1)$. The savings will earn interest at rate $r > 0$, which is random with c.d.f. $F(r)$. In period two, the consumer consumes savings plus interest, so $c_2 = (1 + r)(w - c_1)$. Utility has the additive separable form $u(c_1) + \delta u(c_2)$ where the discount factor δ obeys $0 < \delta < 1$ and u is $\mathcal{C}^2$ with $u' > 0$ and $u'' < 0$.

The consumer's problem is to maximize expected utility by choice of current consumption and hence savings. Utility is then $u(c_1) + \delta u\big((1+r)(w-c_1)\big)$. Expected utility is

$$\begin{aligned} EU(c_1) &= \int \Big[u(c_1) + \delta u\big((1 + r)(w - c_1)\big)\Big]\, dF(r) \\ &= u(c_1) + \delta \int u\big((1 + r)(w - c_1)\big)\, dF(r). \end{aligned}$$

The first-order conditions are

$$u'(c_1) = \delta \int (1 + r)u'\big((1 + r)(w - c_1)\big)\, dF(r)$$

or

$$u'(c_1) = \delta \int (1 + r)u'(c_2)\, dF(r).$$

If there was no uncertainty about the interest rate, the integral would become $(1 + r)u'(c_2)$ and the first-order condition would be $u'(c_1) = \delta(1 + r)u'(c_2)$. ◀

23.3.5 The Market for Insurance: The Problem

Insurance problems are one of the motivations for considering uncertainty. Let's consider the problem of homeowner's insurance.

Example 23.3.4: Homeowner's Insurance. The consumer starts with wealth w, but there is a probability π that he will suffer damage D. Because the consumer is risk averse, the *ex ante* damage goes beyond the expected loss itself. However, insurance is available, which not only covers the loss, but can reduce the loss due to risk aversion.

The consumer can buy \$1 of insurance for a price q. There is no limit concerning how much insurance he can buy other than the budget constraint. For the purpose of computing demand, any $x \geq 0$ is okay.[7]

Let x be the amount of insurance demanded. Purchasing x reduces wealth to $w - xq$, but if the homeowner suffers damage D, the insurance pays x to compensate. Thus wealth is $w - xq$ with probability $(1 - \pi)$ and $w - xq - D + x$ with probability π. Expected wealth is

$$(1 - \pi)(w - xq) + \pi(w - xq + x - D) = w - \pi D + (\pi - q)x.$$

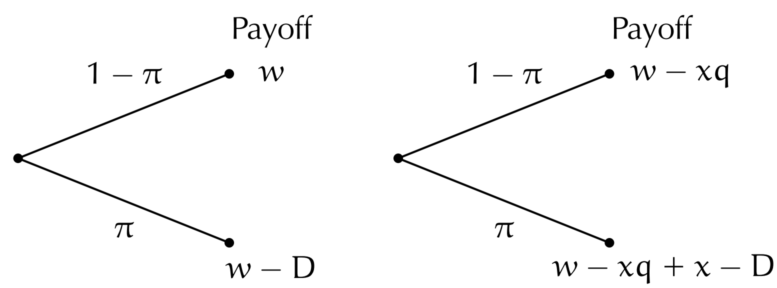

Figure 23.3.5: Payoffs without insurance (left) and payoffs with insurance (right).

[7] Insurance companies usually limit this.

23.3.6 The Market for Insurance: Demand I

Insurance is considered "actuarially fair" if $q = \pi$. In that case expected wealth is $w - \pi D$ regardless of how much insurance the consumer purchases. It is easy to see that due to risk aversion, it will be optimal to purchase exactly enough insurance to offset the loss. By setting $x = D$, the consumer's payoff in either state will be $w - \pi D$. Uncertainty is completely eliminated. Any other option creates a risky lottery with an expected payoff of $w - \pi D$. The risk averse consumer prefers the sure thing!

23.3.7 The Market for Insurance: Demand II

For any other price q, we have to calculate. Expected utility is

$$Eu(x) = (1-\pi)u(w - xq) + \pi u(w - xq + x - D).$$

We must choose x to maximize $Eu(x)$. Assuming an interior solution, the first-order condition is

$$0 = -q(1-\pi)u'(w - xq) + (1-q)\pi u'(w - xq + x - D)$$

We'll deal with the possible corner solutions in a bit. Thus

$$\frac{1-\pi}{\pi}u'(w - xq) = \frac{1-q}{q}u'(w - xq + x - D).$$

If $\pi > q$, $(1-\pi)/\pi < (1-q)/q$ and the first-order conditions yield

$$u'(w - xq) > u'(w - xq + x - D).$$

Diminishing marginal utility then gives

$$w - xq < w - xq + x - D,$$

so $D < x$. The consumer will choose to over-insure in this case where the insurance is underpriced. Note that setting $x = 0$ yields $Eu' > 0$, so the corner solution will not be chosen.

If $\pi < q$, the same argument yields $x < D$. The consumer will under-insure when the price is high and possibly choose the corner solution of no insurance, depending on u and the parameter values.

23.3.8 The Market for Insurance: Supply

We can also ask about the firms offering the insurance. Their expected profit is $xq - x\pi = x(q - \pi)$. Unless they are risk seekers, they will not offer any underpriced insurance. The usual assumption for publicly held firms is that they are approximately risk neutral due to risk spreading and risk pooling on the part of the shareholders. Risk neutrality would lead to unlimited supply at price $q = \pi$. It may not be possible to effectively spread and pool risks that are very costly and affect large areas at the same time (e.g., hurricanes). In that case the price would be expected to exceed π and people will under-insure. ◀

23.3.9 Demand for Risky Assets SKIPPED

We close this series of examples by considering the demand for risky assets.

Example 23.3.6: Asset Demand. We have two assets, a safe asset paying a certain return of $r \geq 1$ and a risky asset paying x which has c.d.f. F. We assume $\mathbf{E}x > r$, otherwise no risk averse consumer would want the risky asset. The consumer starts with wealth w and invests a fraction α in the risky asset and fraction $(1-\alpha)$ in the safe asset. We assume $0 \leq \alpha \leq 1$. The payoff is then $w\big(r + \alpha(x-r)\big)$.

The consumer maximizes expected utility, $\int u\big(wr + w\alpha(x-r)\big)\,dF(x)$ under the constraint $0 \leq \alpha \leq 1$. The derivative is $\Phi(\alpha) = w\int u'\big(wr + w\alpha(x-r)\big)(x-r)\,dF(x)$.

Now $\Phi(0) = w\int u'(wr)(x-r)\,dF = wu'(wr)(\mathbf{E}x - r) > 0$, ruling out $\alpha = 0$ as a solution.

The solution $\alpha = 1$ cannot be immediately ruled out. Here $\Phi(1) = w\int u'(wx)(x - r)\,dF(x)$ which is not easily signed.

For purposes of this example, we specialize to the case where x is uniformly distributed over $[r/3, 2r]$ and the consumer has utility $u(x) = \ln x$. Then

$$F(x) = \begin{cases} 0 & \text{if } x < r/3 \\ \frac{3}{5r}(x - \frac{r}{3}) & \text{if } r/3 \leq x \leq 2r \\ 1 & \text{if } x > 2r \end{cases}.$$

Here $\mathbf{E}F = 7r/6 > r$. Now Φ becomes

$$\Phi(\alpha) = \frac{3w}{5r}\int_{r/3}^{2r} \frac{x-r}{w\big(r+\alpha(x-r)\big)}\,dx = \frac{3}{5\alpha r}\int_{r/3}^{2r}\left[1 - \frac{r}{r+\alpha(x-r)}\right]\,dx.$$

Where we have used the fact that $x - r = (1/\alpha)\big[\big(r + \alpha(x-r)\big) - r\big]$. The integral reduces to

$$\Phi(\alpha) = \frac{1}{\alpha^2}\left[\alpha - \frac{3}{5}\ln\left(\frac{3+3\alpha}{3-2\alpha}\right)\right].$$

We easily calculate $\Phi(1) = [1 - (3/5)\ln 6] < 0$. This rules out the corner solution $\alpha = 1$. Also, because $\Phi(\alpha) > 0$ for small $\alpha > 0$ and $\Phi(1) < 0$, the continuity of Φ implies there is an interior solution via the Intermediate Value Theorem. By examining the derivative of $\alpha^2\Phi$ (which is positive up to $\alpha = 1/2$ and negative thereafter) we find the solution α^* is unique and greater than 1/2. Because the derivative is positive for smaller α and negative for larger α, this is the maximum. The asset-holder will hold more of the risky asset than the safe asset. In this case $\alpha^* \approx 0.738$. ◄

23.4 Arrow-Pratt Risk Aversion

Although we characterized risk aversion five different ways in Theorem 23.1.10, we did not include two popular measures of risk aversion, developed by Pratt (1964) and Arrow (1965). The Arrow-Pratt measures of risk aversion are *absolute risk aversion*

$$r_A(x) = -\frac{u''(x)}{u'(x)}$$

and *relative risk aversion*

$$r_R(x) = -x\frac{u''(x)}{u'(x)}.$$

We have already seen the absolute risk aversion twice, in Lemma 23.1.9 where it is related to the probability premium and in the examples relating the risk premium (23.1.11), risk spreading (23.1.12), and risk pooling 23.1.13.

By definition, $r_R(x) = xr_A(x)$. Both measures are positive when u is concave and negative when u is convex. In other words, positive risk aversion indicates the consumer is risk averse, negative risk aversion indicates the consumer is risk seeking.

23.4.1 Examples of Arrow-Pratt Risk Aversion

Let's calculate risk aversion for a couple of utility functions.

Example 23.4.1: Some Examples. If

$$u(x) = -e^{-ax}$$

with $a > 0$, the absolute risk aversion is constant, $r_A(x) = a$ and the relative risk aversion is increasing, $r_R(x) = ax$.

If

$$u(x) = \frac{1}{1-\sigma}x^{1-\sigma}$$

with $\sigma > 0$, $\sigma \neq 1$, the relative risk aversion is constant, $r_R(x) = \sigma$ and the absolute risk aversion is decreasing $r_A(x) = \sigma/x$. The case $\sigma = 1$ is filled by $u(x) = \ln x$ when $r_R(x) = 1$ and $r_A(x) = 1/x$.

If $u(x) = Cx + D$, both the absolute and relative risk aversion are zero.

These functions and their affine transforms are the only functions with constant absolute or relative risk aversion. We show this in the following proposition. ◀

23.4.2 Constant Risk Aversion

It is fairly easy to characterize utility functions having constant risk aversion, whether absolute or relative.

Proposition 23.4.2. *Suppose the utility function u is $\mathcal{C}^2$ with $u'' < 0$ and $u' > 0$.*

1. *If the absolute risk aversion $r_A(x)$ is constant, then*
$$u(x) = c - be^{-ax}$$
for some real numbers a, b, c with $a, b > 0$.
2. *If the relative risk aversion $r_R(x)$ is constant, then there are real numbers $c, \sigma > 0$ and b with*
$$u(x) = \begin{cases} c + \frac{b}{1-\sigma} x^{1-\sigma} & \textit{for } \sigma \neq 1 \\ c + b \ln x & \textit{for } \sigma = 1. \end{cases}$$

Proof. In the first case,
$$-\frac{u''(x)}{u'(x)} = a$$
for some $a > 0$. Let $v = u'$, and rewrite the differential equation as
$$-\frac{v'}{v} = a \quad \text{or} \quad \frac{dv}{v} = -a\,dx.$$
Integrating, we find $\ln v = -ax + B$ or $v(x) = e^B e^{-ax}$. That is, $u'(x) = e^B e^{-ax}$. Integrating again, we find $u(x) = -(e^B/a)e^{-ax} + c$. Let $b = e^B/a > 0$ to obtain the result.

In the second case,
$$-x\frac{u''(x)}{u'(x)} = \sigma$$
for some $\sigma > 0$. Again set $v = u'$, obtaining
$$-x\frac{v'}{v} = \sigma \quad \text{or} \quad \frac{dv}{v} = -\frac{\sigma}{x}\,dx.$$
Integrating, we find $\ln v = -\sigma \ln x + B$ or $v(x) = e^B x^{-\sigma}$. That is, $u'(x) = bx^{-\sigma}$ where $b = e^B > 0$. Integrating again, we obtain $u(x) = c + bx^{1-\sigma}/(1-\sigma)$ for $\sigma \neq 1$, while if $\sigma = 1$, we find $u(x) = c + b \ln x$. Here $u'' < 0$ requires $\sigma > 0$. □

23.4.3 Six Ways to Think of Risk Aversion 04/5/22

The next three theorems give additional insight concerning the Arrow-Pratt risk aversion and it relation to the other concepts of risk aversion we have examined.

Theorem 23.4.3. *The following are equivalent when u_1, u_2 are $\mathcal{C}^2$ expected utility functions with $u_i' > 0$:*

1. $r_A(x, u_2) \geq r_A(x, u_1)$ *for all* $x > 0$.
2. $r_R(x, u_2) \geq r_R(x, u_1)$ *for all* $x > 0$.
3. *There is a strictly increasing concave function* ψ *with* $u_2 = \psi \circ u_1$.
4. $c(u_2, F) \leq c(u_1, F)$ *for all* F *(equivalently,* $R(x, u_2) \geq R(x, u_1)$*),*
5. $\pi(x, \varepsilon, u_2) \geq \pi(x, \varepsilon, u_1)$ *for all* $x, \varepsilon > 0$.
6. *For* $\bar{x}$ *certain,* $F \succsim_2 \bar{x}$ *implies* $F \succsim_1 \bar{x}$*. That is,* $\int u_2 \, dF \geq u_2(\bar{x})$ *implies* $\int u_1 \, dF \geq u_1(\bar{x})$.

23.4.4 Proof of Theorem 23.4.3

(1) iff (2): Since $r_A = xr_R$, (1) and (2) are equivalent.

(1) implies (3): Since u_1 and u_2 are both increasing functions on $\mathbb{R}$, we can define ψ by $\psi \circ u_1 = u_2$. Apply the chain rule to find $\psi'(u_1)u_1' = u_2'$, implying $\psi' > 0$. Computing the second derivative, we find $\psi''(u_1')^2 + \psi' u_1'' = u_2''$. Using $\psi' = u_2'/u_1'$, we find $\psi''(u_1')^2 = u_2'' - u_2' u_1''/u_1'$. Dividing by u_2' yields $\psi''(u_1')^2/u_2' = r_A(x, u_1) - r_A(x, u_2) \le 0$. It follows that $\psi'' \le 0$ and so ψ is concave.

(3) implies (1): Reversing the argument shows (3) implies (1). Thus (1)-(3) are equivalent.

(3) implies (4): Jensen's inequality shows (3) implies (4) as follows:

$$\begin{aligned} u_2\big(c(u_1, F)\big) &= \psi\Big(u_1\big(c(u_1, F)\big)\Big) = \psi(Eu_1) \\ &\ge E(\psi \circ u_1) = E(u_2) = u_2\big(c(u_2, F)\big). \end{aligned}$$

Since u_2 is increasing, $c(u_2, F) \le c(u_1, F)$.

(4) implies (5): To get from (4) to (5), consider the lottery paying $(x+\varepsilon)$ with probability $(1/2 + \pi)$ and $(x - \varepsilon)$ with probability $(1/2 - \pi)$. Suppose $\pi = \pi(x, \varepsilon, u_1)$. Then $u_1(x) = Eu_1(L)$, so $x = c(u_1, L)$. By (4), $c(u_2, L) \le c(u_1, L) = x$. Then $Eu_2(L) = u_2\big(c(u_2, L)\big) \le u_2(x)$. This means that we may have to increase the probability π above $\pi(x, \varepsilon, u_1)$ to make u_2 indifferent, establishing (5).

(5) implies (1): Now suppose (5) holds. Then

$$\frac{\pi(x, \varepsilon, u_2) - \pi(x, 0, u_2)}{\varepsilon} \ge \frac{\pi(x, \varepsilon, u_1) - \pi(x, 0, u_1)}{\varepsilon}$$

for $\varepsilon > 0$. Take the limit as $\varepsilon \to 0$ and use Lemma 23.1.9 to find $r_A(x, u_2) \ge r_A(x, u_1)$. Thus (1)-(5) are equivalent.

(3) implies (6): Now suppose (1)-(5) hold and $F \succsim_2 \bar{x}$. That means $\int u_2\,dF \ge u_2(\bar{x})$. Now $\int u_2\,dF = \int \psi \circ u_1\,dF \le \psi\big(\int u_1\,dF\big)$ by Jensen's inequality. It follows that $\psi\big(\int u_1\,dF\big) \ge u_2(\bar{x}) = \psi \circ u_1(\bar{x})$. This implies $\int u_1\,dF \ge u_1(\bar{x})$ because ψ is increasing. Thus (6) holds.

(6) implies (4): Finally, suppose (6) holds. We will show (4) is true. Since $F \succsim_2 c(u_2, F)$, $F \succsim_1 c(u_2, F)$. By definition of certainty equivalence, $c(u_1, F) \sim_1 F$. It follows that $c(u_1, F) \succsim_1 c(u_2, F)$. By monotonicity, $c(u_1, F) \ge c(u_2, F)$.

23.4.5 Wealth and Certainty Equivalence: Constant RRA

Let's consider how the certainty equivalent of a lottery is affected by changes in wealth.

Example 23.4.4: Changes in Wealth. We start with a lottery L that pays \$1 with probability 1/2 and \$4 with probability 1/2. We will consider how much owning the lottery adds to wealth, in certainty equivalence terms. Let x be the level of wealth without the lottery. We presume x is certain. We will calculate the certainty equivalent of $x + L$ and see how much it adds to x. In other words, we take $c_{x+L} - x$ as the value of the lottery.

Suppose utility is $u(w) = w^{1/2}$. Here relative risk aversion is constant and absolute risk aversion is decreasing. We will compare $c_L - 0$ and $c_{x+L} - x$. The expected utility from L is $Eu(L) = (1/2)\sqrt{1} + (1/2)\sqrt{4} = 3/2$. The certainty equivalent obeys $u(c_L) = 3/2$, so $c_L = 9/4$. Now consider $x + L$. It yields expected utility $Eu(x + L) = (1/2)(x + 1)^{1/2} + (1/2)(x + 4)^{1/2}$. Then

$$\begin{aligned} c_{x+L} &= \frac{1}{4}(x+1) + \frac{1}{2}(x+1)^{1/2}(x+4)^{1/2} + \frac{1}{4}(x+4) \\ &= \frac{x}{2} + \frac{5}{4} + \frac{1}{2}(x+1)^{1/2}(x+4)^{1/2}. \end{aligned}$$

We obtain

$$c_{x+L} - x = -\frac{x}{2} + \frac{5}{4} + \frac{1}{2}(x+1)^{1/2}(x+4)^{1/2}.$$

It follows that the change in the value of the certainty equivalent is

$$c_{x+L} - x - c_L = -\frac{x}{2} - 1 + \frac{1}{2}(x+1)^{1/2}(x+4)^{1/2} > 0.$$

This will hold if and only if $(1/2)(x + 1)^{1/2}(x + 4)^{1/2} > 1 + x/2$. Square both sides to see that it holds. Thus $c_{x+L} - x > c_L - 0$.

The certainty equivalent increases faster than wealth because of the decreasing absolute risk aversion.

23.4.6 Wealth and Certainty Equivalence: Constant ARA

Now consider the constant absolute risk aversion case where $u(w) = -e^{-w}$. Then

$$Eu(L) = -(1/2)\left[e^{-1} + e^{-4}\right]$$

so

$$c_L = -\ln\left\{\frac{1}{2}[e^{-1} + e^{-4}]\right\}.$$

Also,

$$Eu(x + L) = -\frac{1}{2}\left[e^{-x-1} + e^{-x-4}\right] = e^{-x}Eu(L),$$

so $c_{x+L} = -\ln e^{-x} + c_L = x + c_L$. In this case, $c_{x+L} - x = c_L$ is independent of the wealth level x.

This illustrates the economic significance of risk aversion changes as wealth changes. Declining absolute risk aversion was paired with an increasing value of the lottery (with wealth), while constant absolute risk aversion led to a constant value of the lottery (regardless of wealth). Example 23.4.6 explores this point further. ◀

23.4.7 Absolute Risk Aversion and Wealth

In fact, changes in absolute risk aversion relate to changes in the absolute level of wealth while changes in relative risk aversion are related to changes in relative wealth. The details are spelled out in the next two theorems.

Theorem 23.4.5. *The following are equivalent when u is a $\mathcal{C}^2$ expected utility function with $u' > 0$:*

1. *Absolute risk aversion is a weakly decreasing function.*
2. *When $x_2 < x_1$, the function $u_2(z) = u(x_2 + z)$ is a concave transformation of $u_1(z) = u(x_1 + z)$.*
3. *For any c.d.f. $F(z)$, the certainty equivalent of adding risk z to wealth x (denoted c_x and defined by $u(c_x) = \int u(x + z)\, dF(z)$) has the property that $x - c_x$ is weakly decreasing in x.*
4. *$\pi(x, \varepsilon, u)$ is weakly decreasing in x.*
5. *For any c.d.f. $F(z)$, $\int u(x_2 + z)\, dF(z) \geq u(x_2)$ and $x_2 < x_1$ imply that $\int u(x_1 + z)\, dF(z) \geq u(x_1)$.*

Proof. This follows easily from Theorem 23.4.3. E.g., If absolute risk aversion is weakly decreasing, then $r_A(x, u_2) \geq r_A(x, u_1)$ which yields (2) by theorem 23.4.3. The other parts follow easily. $\square$

23.4.8 More on Wealth and Certainty Equivalence

SKIPPED REST OF CHAPTER

We can use the properties of absolute risk aversion to examine the relation between wealth and risk aversion.

Example 23.4.6: Wealth and Risk Aversion. Start with wealth w. Lottery L offers a payoff of G (good!) with probability p and B (bad!) with probability $(1-p)$. Here $G > B$. We consider two scenarios. In the first, the consumer owns L in addition to his wealth. At what price will the consumer be willing to sell the lottery?

To sell at price s requires $u(w+s) \geq EU(w+L)$. We find the minimum acceptable sales price by solving $u(w+s) = EU(w+L)$. Using the notation of the previous theorem, $w+s = c_{w+L}$ or $s = c_{w+L} - w$.

In the second scenario, the consumer does not own L. At what price q will the consumer be willing to buy L? The consumer is willing to buy if $Eu(w-q+L) \geq u(w)$. The maximum price when the consumer is willing to buy solves $Eu(w-q+L) = u(w)$. Thus $c_{w-q+L} = w$, equivalently, $q = c_{w-q+L} - (w-q)$.

If absolute risk aversion is decreasing, we know $w - c_{w+L} < (w-q) - c_{w-q+L}$, so $-s < -q$ or $q < s$. If absolute risk aversion is increasing, $q > s$. If the absolute risk aversion is constant, $q = s$. This makes sense if we think about the wealth differences in the two cases. When risk aversion is decreasing, risk is less costly when wealthy, meaning that the lottery will trade at a higher price. Thus $s > q$. If risk aversion is increasing in wealth, the cost of the lottery's risk will be higher in the wealthy case, causing the lottery's price to be lower, $s < q$. ◄

23.4.9 Relative Risk Aversion and Wealth

Theorem 23.4.7. *The following are equivalent when u is a $\mathcal{C}^2$ expected utility function with $u' > 0$:*

1. *Relative risk aversion is a weakly decreasing function.*
2. *When $x_2 < x_1$, the function $u_2(t) = u(tx_2)$ is a concave transformation of $u_1(t) = u(tx_1)$.*
3. *For any c.d.f. $F(z)$, the certainty equivalent of multiplying wealth x by a risk z (denoted $\bar{c}_x$ and defined by $u(\bar{c}_x) = \int u(zx)\,dF(z)$) has the property that $x/\bar{c}_x$ is weakly decreasing in x.*

Proof. The proof is similar to Theorem 23.4.5. For example,

$$r_R(t, u_i) = -\frac{u''(tx_i)tx_i^2}{u'(tx_i)x_i} = -\frac{u''(tx_i)tx_i}{u'(tx_i)}.$$

When $x_2 < x_1$ and relative risk aversion is weakly decreasing (1), $r_R(t, u_1) \leq r_R(t, u_2)$. Then (2) holds by Theorem 23.4.3. □

23.5 Stochastic Dominance

Stochastic dominance is a way to compare risks that is independent of the utility function. There are several types of stochastic dominance.

First-order stochastic dominance. Given two c.d.f.'s F and G, F is *first-order stochastically dominant* over G if for all weakly increasing $u\colon \mathbb{R} \to \mathbb{R}$, we have $\int u(x)\,dF(x) \geq \int u(x)\,dG(x)$.

First-order stochastic dominance necessarily affects the mean. Setting $u(x) = x$ shows that if F stochastically dominates G, then the mean of F is at least as large as the mean of G. First-order dominance does not necessarily involve an increase in risk, but may only indicate an increase in the mean.

Recall that $E(X^2) = \sigma^2 + \mu^2$ where μ is the mean and σ is the standard deviation. When F first-order dominates G, we can set $u(x) = x^2$ to find $\sigma_F^2 + \mu_F^2 \geq \sigma_G^2 + \mu_G^2$, but if $\mu_F > \mu_G$, we cannot conclude that $\sigma_F \geq \sigma_G$.

Stochastic dominance applies when we shift the distribution to the left or right. Let F be a c.d.f., and consider $F_z(x) = F(x + z)$. We have just shifted the distribution over. The only moment that has changed is the mean. The variance and higher moments are unchanged as both the mean and distribution function are shifted by the same amount.

If μ is the mean of F, then $\mu - z$ is the mean of F_z. For $z > 0$, the mean is reduced and F stochastically dominates F_z as $F_z(x) \leq F(x)$. For $z < 0$, the reverse holds.

23.5.1 First-Order Stochastic Dominance

The next result is that F first-order dominates G if and only if F's distribution is to the right of G's distribution.

Proposition 23.5.1. *A c.d.f.* F *first-order dominates* G *if and only if* $F \le G$.

Proof. **Part I**: Suppose F first-order dominates G and let $H(x) = F(x) - G(x)$. If there is a y such that $H(y) > 0$, let

$$u(x) = \begin{cases} 1 & \text{when } x \ge y \\ 0 & \text{when } x < y. \end{cases}$$

Then $\int_{-\infty}^{+\infty} u\, dH = \int_{y}^{+\infty} dH = -H(y) < 0$, contradicting $\int u\, dH = \int u\, dF - \int u\, dG \ge 0$. Thus $H(x) \le 0$ and so $F(x) \le G(x)$ for all x.

Part II: Suppose $F \le G$ and define $H = F - G \le 0$. It is enough to show this for all weakly increasing step functions u.[8] A step function is defined by $x_1 < x_2 < \cdots < x_N$ and α_i, $i = 0, \dots, N$. We have $u(x) = \alpha_0$ when $x < x_1$, $u(x) = \alpha_i$ for $x_i \le x < x_{i+1}$ and $i = 1, \dots, N-1$ and $u(x) = \alpha_N$ for $x_N \le x$. That u is weakly increasing means $\alpha_i \le \alpha_{i+1}$. We can presume $\alpha_i < \alpha_{i+1}$ without loss of generality (by combining steps if needed).

Now

$$\begin{aligned} \int u\, dH &= \alpha_0 H(x_1) + \alpha_1\big(H(x_2) - H(x_1)\big) + \cdots \\ &\qquad + \alpha_{N-1}\big(H(x_N) - H(x_{N-1})\big) - \alpha_N H(x_N) \\ &= (\alpha_0 - \alpha_1)H(x_1) + \cdots + (\alpha_{N-1} - \alpha_N)H(x_N) \\ &\ge 0 \end{aligned}$$

since both $H \le 0$ and $\alpha_i - \alpha_{i+1} \le 0$. This implies $\int u\, dF - \int u\, dG = \int u\, dH \ge 0$ for all weakly increasing u, which establishes stochastic dominance. □

[8] This is a standard result from measure theory.

23.5.2 Stochastic Dominance with Finitely Many States

Stochastic dominance applies to discrete lotteries as well as continuous lotteries. Let L and L′ be lotteries on $\{1, 2, \ldots, N\}$ with probabilities p_n and p'_n. Then L first-order stochastically dominates L′ if $\Pr(L \leq k) \geq \Pr(L' \leq k)$. That is, if

$$\sum_{n=1}^{k} p_n \geq \sum_{n=1}^{k} p'_n$$

for all k.

23.5.3 Second-Order Stochastic Dominance

Second-order stochastic dominance allows us to compare the riskiness of two random variables with the same mean.

Second-order Stochastic Dominance. Given two c.d.f.'s F and G, F is *second-order stochastically dominant* over G if

1. F and G have the same mean.
2. For all weakly increasing and concave $u\colon \mathbb{R} \to \mathbb{R}$, we have

$$\int u(x)\, dF(x) \geq \int u(x)\, dG(x).$$

In other words, if u is a concave utility function, the utility of someone who is risk averse, then $Eu(F) \geq Eu(G)$ whenever F second-order dominates G. Risk averse consumers prefer distributions that are second-order stochastically dominant.

23.5.4 Mean-Preserving Spreads Increase Risk

One way to increase risk in a stochastic dominance sense is via a mean-preserving spread.

Example 23.5.2: Continuous Mean-preserving Spread. We start with a lottery $F(x)$. We will add some noise to F without affecting the mean. Suppose $H_x(z)$ be a family of c.d.f.'s that has mean zero for every x. Then $x + z$ is a mean-preserving spread of x. The mean of $x + z$ is $\int \left(\int (x + z)\, dH_x(z) \right) dF(x) = \int x\, dF(x)$, so the mean of x and $x + z$ are the same. Let G be the corresponding reduced lottery and suppose u is weakly increasing and concave.

$$\begin{aligned}
\int u(x)\, dG(x) &= \int \left(\int u(x + z)\, dH_x(z) \right) dF(x) \\
&\leq \int u \left(\int (x + z)\, dH_x(z) \right) dF(x) \\
&= \int u(x)\, dF(x)
\end{aligned}$$

which establishes second-order stochastic dominance of F over the noisier G. ◄

23.5.5 Risk and Asset Demand I

We will use a mean-preserving spread to see how risk affects asset demand.

Example 23.5.3: Risk and Asset Demand. This is based on Example 23.3.6. Recall that we have two assets, one risky and one safe. We make the slight change that the safe asset pays a return of 1. The risky asset's return x has distribution function $F(x)$. We assume that $Ex = r > 1$, $x > 0$ (no negative returns) and that $x < 1$ with positive probability. Let α be the share of wealth invested in the risky asset, with the remainder invested in the safe asset. The derivative of expected utility with respect to α is

$$\phi(\alpha) = w \int u'\big(w + w\alpha(x-1)\big)(x-1)\,dF(x).$$

Define ψ by $\psi(x,\alpha) = u'\big(w + w\alpha(x-1)\big)(x-1)$, so $\phi(\alpha) = w\int \psi(x,\alpha)\,dF(x)$. Now $\psi_\alpha = wu''(x-1)^2 < 0$, so $\phi' < 0$. This implies there is at most one maximizing α and that it is the only solution to the first-order equation $\phi(\alpha) = 0$. As $\phi(0) = u'(w)(r-1) > 0$, the solution $\alpha^* > 0$.

23.5.6 Risk and Asset Demand II

Consider the mean-preserving spread H_x and define

$$\tilde{\phi}(\alpha) = w \int \left(\int \psi(x+z, \alpha)\, dH_x(z) \right)\, dF(x).$$

If ψ is concave, Jensen's inequality implies

$$\int \psi(x+z, \alpha)\, dH_x(z) \le \psi \left(\int (x+z)\, dH_x(z) \right) = \psi(x, \alpha).$$

Thus $\tilde{\phi}(\alpha) \le \phi(\alpha)$. The downward shift in ϕ results in a lower value of α^*, the demand for the risky asset. If ψ is convex, the inequality reverses, ϕ shifts up, and demand for the risky asset increases.

When is ψ concave? We have $\psi_{xx} = w\alpha[2u'' + w\alpha u'''(x-1)]$ with both u'' and u''' evaluated at $c = w + w\alpha(x-1)$. It follows that any risk averse utility function with $u''' < 0$ will behave in the expected way, with an increase in risk reducing demand for the risky asset. If $u''' > 0$, there is a possibility that an increase in risk increases demand for the risky asset.

We now consider the case $u(c) = c^{1-\sigma}/(1-\sigma)$ (or $u(c) = \ln c$ when $\sigma = 1$).[9] Here

$$\psi_{xx}(x, \alpha) = \frac{w^2 \alpha \sigma}{\big(w + w\alpha(x-1)\big)^{\sigma+2}}[-2 + \alpha(\sigma - 1)(x-1)].$$

In the logarithmic case ($\sigma = 1$), $\psi'' < 0$ and we find that asset demand decreases as risk increases. In fact, this will be true ◄

[9] These functions are well-defined because we have assumed $x > 0$.

23.6 The Limits of Expected Utility

We can see what special preferences expected utility represents by drawing indifference curves in the Marschak-Machina triangle. Indifference curves for expected utility are always parallel straight lines. The common slope will be different for different utility functions, but the indifference curves are always parallel straight lines.

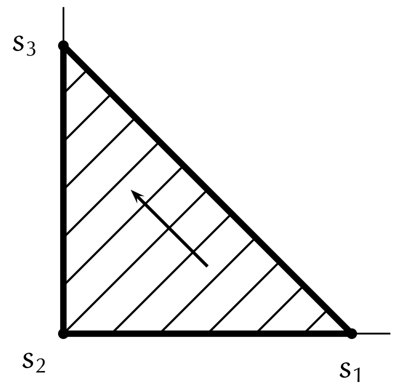

Figure 23.6.1: Typical Marschak-Machina Triangle. There are 3 states with $s_1 \prec s_2 \prec s_3$. The indifference curves are parallel straight lines for every expected utility function. Their common slope is $-[u(s_1) - u(s_2)]/[u(s_3) - u(s_2)]$

The response to uncertainty is what's being captured in the Marschak-Machina triangle. The fact that the utility function itself may non-linear allows substantial flexibility, but the treatment of uncertainty is heavily constrained, yielding parallel indifference curves. In spite of this restriction, expected utility proven very fruitful in the analysis of the economics of uncertainty in all manner of economic problems. In particular, it is a workhorse when dealing with the economics of financial markets.

Nonetheless, this places a substantial limitation on preferences, a limitation that sometimes makes it impossible to model real-world decision making without greatly complicating our models.

In this section we consider issues that involve a failure of the basic axioms. We will consider this in the context of von Neumann-Morgenstern utility, where the Independence Axiom may fail.[10]

[10] In the decision theory context of Chapter 24, this generally translates into a failure of the Sure Thing Principle.

23.6.1 The Limits of Expected Utility

Although expected utility has proven to be a powerful technique for studying decisions under uncertainty, it does impose strong restrictions on preferences. Indeed, some consider the restrictions to be contrary to intuition. Because of this, alternatives to the Independence Axiom are sometimes considered.

In some cases expected utility behaves contrary to intuition. Some of these cases are mistakes. Once people fully understand the lottery they face, they do not want to violate expected utility. Other such situations do not appear to be mistakes. They either involve issues that are not addressed by the expected utility setup (e.g., the timing of resolution of uncertainty, as handled by Kreps-Porteus preferences, Kreps and Porteus, 1978) or other issues that we don't fully understand.

23.6.2 Allais Paradox

The following example will help clarify the restrictions imposed by the Independence Axiom.

Example 23.6.2: Allais Paradox. Suppose $S = 3$. You face a lottery with 3 possible prizes: \$5 million, \$1 million, and \$0. Lottery 1 pays \$1 million with probability 1 while lottery 2 pays \$5 million with probability 0.1, \$1 million with probability 0.89 and \$0 with probability 0.01. Most people prefer lottery 1, with its certain large gain, to the risk of losing everything in lottery 2.

Lottery 3 pays \$1 million 11% of the time and \$0 89% of the time. Lottery 4 pays \$5 million 10% of the time and \$0 90% of the time. Most people prefer lottery 4 to lottery 3.[11]

These choices violate expected utility. Let u_2 be the utility of \$2.5 million, u_1 the utility of \$500,000, and u_0 the utility of \$0. Since lottery 1 is preferred to lottery 2, we find $u_1 > 0.1u_2 + 0.89u_1 + 0.01u_0$. It follows that $0.11u_1 > 0.1u_2 + 0.01u_0$. We also have lottery 4 preferred to lottery 3. That means $0.11u_1 + 0.89u_0 < 0.1u_2 + 0.9u_0$. Rearranging, we find $0.11u_1 < 0.1u_2 + 0.01u_0$, contradicting the results from the first pair of lotteries. Thus these preferences are inconsistent with expected utility. See Allais (1953). ◀

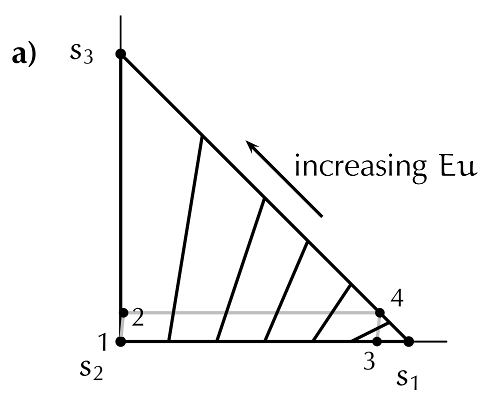

Figure 23.6.3: Here s_1 is a payoff of \$0, s_2 is a payoff of \$1 million, and s_3 pays \$5 million. The Allais lotteries are denoted 1–4. Thus 1–4 forms a parallelogram with the lines connecting 1–2 and 3–4 both having the same slope. When expected utility indifference curves are parallel straight lines, $L_1 \succ L_2$ if and only if $L_3 \succ L_4$. To get the Allias result, the slopes must increase as utility increases, causing the indifference curves to fan out as shown here.

[11] See Allais, 1953; Morrison, 1967; Raiffa, 1968; Slovic and Tversky, 1974.

23.6.3 Allais and the Common Consequence Effect

To understand the Allais paradox a bit better, recall that the lotteries are:

$$
\begin{aligned}
L_1 &= \delta_1 & L_2 &= 0.1\delta_5 \oplus 0.89\delta_1 \oplus 0.01\delta_0 \\
L_3 &= 0.11\delta_1 \oplus 0.89\delta_0 & L_4 &= 0.1\delta_5 \oplus 0.9\delta_0
\end{aligned}
$$

where the amounts x in each δ_x are in millions. Now define $L = \frac{10}{11}\delta_5 \oplus \frac{1}{11}\delta_0$, $L' = \delta_0$ and $L'' = \delta_1$.

The Allais paradox is not the only case where there's evidence that the indifference curves get steeper. There are a whole set of such cases. We can expose the structure of the Allais paradox by rewriting the lotteries as

$$
\begin{aligned}
L_1 &= \alpha\delta_x \oplus (1-\alpha)L'' & L_2 &= \alpha L \oplus (1-\alpha)L'' \\
L_3 &= \alpha\delta_x \oplus (1-\alpha)L' & L_4 &= \alpha L \oplus (1-\alpha)L'
\end{aligned}
\tag{23.6.3}
$$

where $x = 1$ and $\alpha = 0.11$. In the Allais case L involves payoffs both greater and less that $x = 1$ and that L'' stochastically dominates L'.

Consider the pattern above for general x, α and lotteries L, L', and L''. In each row of equation 23.6.6, we mix the same lotteries with a common third lottery, either L' or L''. If $\delta_x \succsim L$, the Independence Axiom implies $L_1 \succsim L_2$ and $L_3 \succsim L_4$ while if $\delta_x \prec L$, Independence implies $L_2 \succsim L_1$ and $L_4 \succsim L_3$. But then, when preferences obey the von Neumann-Morgenstern Axioms, $L_1 \succsim L_2$ if and only if $L_3 \succsim L_4$. The preferences must run in the same direction when mixing with the same lottery.

Nonetheless, people often choose L_1 and L_4 in lotteries of this form.[12] This type of preference reversal is referred to as the *common consequence effect*.. The *common consequence effect* occurs whenever mixture of two options with a common third option can reverse preferences. Here it happens when L'' is used to mix rather than L'. If the lotteries involve only three states, we again have indifference curves fanning out whenever there is a common consequence effect.

[12] See MacCrimmon, 1968; MacCrimmon and Larsson, 1979; Kahneman and Tversky, 1979; Chew and Waller, 1986.

23.6.4 The Common Ratio Effect

The Allais paradox is not the only type of choice where people choose in a way that violates the von Neumann-Morgenstern Axioms. In its simplest form, we consider lotteries involving zero and each of two certain outcomes x and y. Then mix the same fraction r of these lotteries with zero.

$$\begin{array}{ll} L_1 = \alpha\delta_x \oplus (1-\alpha)\delta_0 & L_2 = \beta\delta_y \oplus (1-\beta)\delta_0 \\ L_3 = r\alpha\delta_x \oplus (1-r\alpha)\delta_0 & L_4 = r\beta\delta_y \oplus (1-r\beta)\delta_0. \end{array} \tag{23.6.4}$$

Once again, the von Neumann-Morgenstern Axioms imply that $L_1 \succsim L_2$ if and only if $L_3 \succsim L_4$. Nonetheless, in practice this often leads to a preference reversal, called the *common ratio effect*. Here the probabilities of x and y have been reduced by the same factor, leaving the relative probability of x and y unchanged.

The common ratio effect also involves a fanning of indifference curves. It lies behind a couple of other known paradoxes. When $\alpha = 1$, this is the "certainty effect" of Kahneman and Tversky (1979). By setting $y = 2x$ and $\alpha = 2\beta$ we can also obtain the "Bergin Paradox" of Hagen (1979).[13] Yaari's (1987) dual theory of utility avoids both of these problems, but at the cost of creating a different set of problems.

[13] See Machina (1987).

23.6.5 Edwards' Framework

Other attempts have been made to deal with these issues. Edwards (1954, 1955) used a model where utility is given by

$$\sum_{s=1}^{S} \pi(p_s)u(x_s)$$

where x_s is the payoff in state s and $\pi(p_s)$ is a subjective "probability" or "decision-weight" derived from the objective probability. Edwards noted that the probability transformation π need not be additive. This implies that $\sum_s \pi(p_s)$ need not be one.

Kahneman and Tversky (1979) used Edwards' framework when formulating prospect theory, which can also deal with the Allais paradox. They considered choices over lotteries, as in the Von Neumann-Morgenstern framework. However, these lotteries are then considered with reference to the status quo, and undergo a mental editing process to simplify choice. Finally, they apply Edwards' framework of decision weights with a function u that is concave for gains and convex for losses, with $u'(0-)$ being its steepest slope. The reference point used has a big impact on the resulting utility.[14]

These and other alternative utility specifications has been tried to deal with the problems of expected utility, but they have not been generally adopted.

April 6, 2022

[14] Markowitz (1952b) proposed using current wealth as a reference point in his discussion of the Friedman and Savage's (1948) utility function.