Money Lotteries and Risk Aversion

Uncertainty arises in economic settings. For example, a consumer might be deciding how much of her savings to put into a safe asset and how much to put into a risky asset.

Settings like this might require us to consider an infinite number of alternatives and an infinite number of outcomes.

For simplicity we will consider a single consumption good, so outcomes are in terms of money and alternatives are lotteries over money.

Def. A money lottery is a cumulative distribution function, $F : \Re \rightarrow [0, 1]$. That is, for any amount of money x, $F(x)$ is the probability that the lottery pays less than or equal to x.

If the cdf has an associated density function f and is defined over the support $[\underline{x}, \overline{x}]$, then we have $F(x) = \int_{\underline{x}}^{x} f(t)dt$.

For a compound lottery $(F_1, ..., F_K; \alpha_1, ..., \alpha_K)$, we can convert it into a simple lottery by computing the cdf based on the cdf's for the underlying simple lotteries:

$$F(x) = \sum_{k=1}^{K} \alpha_k F_k(x).$$

We now take the space of lotteries $\mathcal{L}$ to be the space of distribution functions over nonnegative amounts of money, although specific environments may impose further restrictions on the support.

The expected utility theorem (for continuous outcomes) says that, under continuity and independence, there is a utility function over outcomes, $u(x)$, such that the utility of any lottery is given by

$$U(F) = \int u(x) dF(x).$$

$U(\cdot)$ is the v.N-M expected utility function and $u(\cdot)$ is called the Bernoulli utility function. We assume that $u(\cdot)$ is strictly increasing (monotonic) and continuous.

Def. A DM is *risk averse* if for any lottery F, the degenerate lottery yielding the certain outcome $\int x dF(x)$ is weakly preferred to the lottery F. If the DM is always indifferent between the two lotteries, she is *risk neutral*, and if $\int x dF(x)$ with certainty is always strictly preferred (unless F is also degenerate) then the DM is *strictly risk averse*.

Based on this definition, a v.N-M DM is risk averse if and only if

$\int u(x) dF(x) \leq u(\int x dF(x))$ holds for all lotteries $F(\cdot)$.

The above inequality is called Jensen's inequality, and it is true if and only if u is a concave function.

Thus, risk aversion is equivalent to the Bernouilli utility function being concave. Risk neutrality is equivalent to u being linear.

The word "aversion" makes sense, because the DM would not be willing to take a fair bet. (Draw a graph.)

Def. Given a Bernoulli utility function u,

(i) The *certainty equivalent* of F, denoted by $c(F, u)$, is the amount of money for which the DM is indifferent between the lottery F and the certain amount c solving

$$u(c) = \int u(x) dF(x).$$

(ii) For any fixed amount of money x and positive "bet" amount ε, the *probability premium*, denoted by $\pi(x, \varepsilon, u)$ is the excess in winning probability (above the fair odds of 0.5) that makes the DM indifferent between the certain outcome x and the gamble between the outcomes $x + \varepsilon$ and $x - \varepsilon$. That is, $\pi(x, \varepsilon, u)$ solves

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

Proposition: Suppose a DM is an expected utility maximizer with a Bernoulli utility function over money, $u(\cdot)$. Then the following properties are equivalent:

(i) the DM is risk averse,

(ii) $u(\cdot)$ is concave,

(iii) $c(F, u) \leq \int x dF(x)$ for all $F(\cdot)$,

(iv) $\pi(x, \varepsilon, u) \geq 0$ for all x, ε.

Example (Insurance) DM has initial wealth W. Loses D w.p. π and loses nothing w.p. $1-\pi$. The Bernoulli utility function is the natural logarithm, $u(x) = \log(x)$.

The price per unit of insurance is q. (if you buy α units, you pay $q\alpha$, and if you have a loss you collect α)

What is the optimal α?

$$\max_{\alpha} (1-\pi)\log(W - q\alpha) + \pi \log(W - q\alpha - D + \alpha)$$

If the price of insurance is close enough to "fair odds," then there is an interior solution solving the f.o.c.

$$\frac{q(1-\pi)}{\pi(1-q)} = \frac{W - \alpha q}{W - D + \alpha(1-q)}.$$

Solving for α, we have

$$\alpha^* = \frac{\pi W}{q} - \frac{(W-D)(1-\pi)}{1-q}.$$

If $\alpha^* < 0$, then the DM is at a corner in which she strictly prefers not to buy insurance.

Under fair odds, the insurance payment per unit, q, equals the expected claim per unit, π. Then $\alpha^* = D$. The DM eliminates all risk and consumes the expectation of her after-loss wealth, $W - \pi D$.

If $q > \pi$, the DM chooses less than full insurance.

What is the certainty equivalent of the original lottery with no insurance? Solve

$$\begin{aligned}\log(c) &= \int \log(x) dF(x) \\ &= (1-\pi)\log(W) + \pi \log(W - D). \\ c &= W^{1-\pi}(W-D)^{\pi}.\end{aligned}$$

For example, if $W = 1$, $D = \frac{3}{4}$, and $\pi = \frac{1}{2}$, the expection of the lottery is $\frac{5}{8}$ but the certainty equivalent is $\frac{1}{2}$.

With log utility, $x = 1$, and $\varepsilon = \frac{1}{2}$, what is the probability premium? Solve

$$\begin{aligned}
\log(1) &= (\frac{1}{2} + \pi)\log(\frac{3}{2}) + (\frac{1}{2} - \pi)\log(\frac{1}{2}) \\
0 &= \frac{1}{2}\log(\frac{3}{4}) + \pi\log(3) \\
\pi &= -\frac{\frac{1}{2}\log(\frac{3}{4})}{\log(3)} \simeq 0.1309
\end{aligned}$$

Example (A Portfolio Problem) The DM has initial wealth W and must decide how much to allocate to each of two assets.

Asset 1 pays a safe return, so x^1 units invested yields x^1 units of consumption.

Asset 2 is a risky asset paying a higher expected return. x^2 units invested yields Rx^2 units of consumption, where $R = 3$ with probability $\frac{1}{2}$ and $R = 0$ with probability $\frac{1}{2}$.

If the Bernoulli utility function over final consumption is $u(x) = x^{1/2}$, what is the optimal portfolio?

$$\begin{aligned} & \max \frac{1}{2}(x^1 + 3x^2)^{1/2} + \frac{1}{2}(x^1 + 0x^2)^{1/2} \\ & \text{subject to} \\ x^1 + x^2 &= W \end{aligned}$$

Solving by either the Lagrangean method or substitution, we have

$$x^1 = \frac{W}{2} \quad \text{and} \quad x^1 = \frac{W}{2}.$$

Final consumption is $2W$ when the risky asset return is high and $\frac{W}{2}$ when the risky asset return is zero.

Measures of Risk Aversion

Def. The *Arrow-Pratt coefficient of absolute risk aversion* at x is

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

$u''(x)$ is a measure of the curvature of the utility function, but is affected by linear transformations. This is why we divide by $u'(x)$.

A risk neutral person will have $r_A(x) = 0$.

The exponential utility function $u(x) = -e^{-ax}$ has a constant coefficient of absolute risk aversion:

$$-\frac{u''(x)}{u'(x)} = -\frac{-a^2 e^{-ax}}{a e^{-ax}} = a$$

Def. The *coefficient of relative risk aversion* at x is

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

Since $r_R(x) = xr_A(x)$ holds, it is clear that decreasing relative risk aversion implies decreasing absolute risk aversion. Why?

The CRRA utility function $u(x) = \frac{x^{1-a}}{1-a}$ (for $a \geq 0$) exhibits constant relative risk aversion:

$$-\frac{xu''(x)}{u'(x)} = -\frac{x(-ax^{-a-1})}{x^{-a}} = a.$$

Constant absolute risk aversion means that the amount of money the DM is willing to pay to eliminate a fair bet of \$100 does not depend on her wealth. Constant relative risk aversion means that the percentage of wealth the DM is willing to pay to eliminate a fair bet of 10% of her wealth does not depend on her wealth.

Stochastic Dominance

When does one lottery unambiguously yield higher utility than another?

When is one lottery unambiguously less risky than another?

Definition: A money lottery F f*irst-order stochastically dominates* (fosd) a money lottery G if for every non-decreasing function, $u : \Re \rightarrow \Re$, we have

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

Interpretation of fosd: every expected utility maximizer, weakly preferring more money to less money, prefers the lottery F to G.

Proposition: F fosd G if and only if $F(x) \leq G(x)$ holds for all x.

Proof of "only if": Define $H(x) = F(x) - G(x)$. Suppose $F(x) \leq G(x)$ does not hold for some x. Then there is x^* such that $H(x^*) > 0$.

Define the non-decreasing function u by: $u(x) = 1$ for $x > x^*$ and $u(x) = 0$ for $x \leq x^*$. Then

$$\begin{aligned} \int_0^\infty u(x) dH(x) &= \int_{x^*}^\infty dH(x) = H(\infty) - H(x^*) \\ &= -H(x^*) < 0. \end{aligned}$$

Thus,

$$\int_0^\infty u(x) dF(x) - \int_0^\infty u(x) dG(x) < 0$$

holds, contradicting that F fosd G.

Proof of "if" (i.e., $F(x) \leq G(x)$ holds for all x implies F fosd G):

We will prove this under the additional assumption of differentiable u. Integration by parts yields

$$\int_0^\infty u(x)dH(x) = u(x)H(x)|_0^\infty - \int_0^\infty u'(x)H(x)dx.$$

Since $F(0) = G(0) = 0$ and $F(\infty) = G(\infty) = 1$, we have

$$\int_0^\infty u(x)dH(x) = -\int_0^\infty u'(x)H(x)dx.$$

Since $H(x) \leq 0$ holds by assumption, and $u'(x) \geq 0$, the right side is non-negative, so we have $\int_0^\infty u(x)dF(x) \geq \int_0^\infty u(x)dG(x)$. That is, F fosd G.

Definition: For two lotteries, F and G, with the same mean, F *second-order stochastically dominates* G (sosd, or, F is less risky than G) if for every non-decreasing concave function $u : \Re \to \Re$, we have

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

Interpretation of sosd: every risk averse expected utility maximizer prefers the lottery F to G.

Definition: G is a *mean-preserving spread* of F if there exist distributions $H_x(z)$ such that (i) $H_x(z)$ has mean zero for all x, and (ii) $G(x) = pr(\overline{x} + z \leq x)$, where $\overline{x}$ is the outcome of lottery F and z is the outcome of lottery $H_{\overline{x}}(z)$.

Interpretation of mean-preserving spread: G is a compound lottery, where first we draw an outcome of F, then draw an outcome of H, then add the two outcomes. The second step does not change the mean but "adds noise."

Illustrate with an example of "elementary increase in risk."

Notice that if G is a mean-preserving spread of F and u is concave, then F sosd G. To see this,

$$
\begin{aligned}
\int u(x)dG(x) &= \int_{\overline{x}} \left[\int_z u(\overline{x}+z)dH_{\overline{x}}(z) \right] dF(\overline{x}) \\
&\leq \int_{\overline{x}} u(\int_z (\overline{x}+z)dH_{\overline{x}}(z)dF(\overline{x}) \\
&= \int_{\overline{x}} u(\overline{x})dF(\overline{x}) = \int u(x)dF(x).
\end{aligned}
$$

In fact, we can show:

Proposition: If F and G have the same mean, then the following are equivalent:

(i) F sosd G,

(ii) G is a mean-preserving spread of F,

(iii) for all x, we have

$$\int_0^x G(t)dt \geq \int_0^x F(t)dt.$$

Representing Uncertainty by States of Nature

The outcomes associated with a lottery are generated by some underlying cause: whether you have a car accident for the insurance example, whether the risky firm's invention succeeds in the portfolio example.

When we consider market economies with uncertainty, we must keep track of these underlying causes in order to know whether everyone's lottery over consumption can be implemented at the same time. We call these causes states or *states of nature*. For example, the number of oranges that can be consumed in total depends on the state of nature.

Knowing the state resolves all uncertainty faced by the DM, or more generally, faced by anyone in the economy.

We denote the set of states by S and an individual state by $s \in S$.

The (objective) probability of state s is π_s.

Continuing to assume that there is only one commodity (money), a *random variable* representing state-contingent consumption is a function $g : S \rightarrow \Re_+$ that maps states into monetary outcomes.

Note that g also defines a money lottery F, with $F(x) = \sum_{s:g(s) \leq x} \pi_s$.

The random variable g can be represented by the vector of state-contingent consumptions, $(x_1, ..., x_S)$.

The v.N-M expected utility function is given by

$$\sum_{s=1}^{S} \pi_s u(x_s).$$

For some applications, the utility function itself may depend on the state. For example, maybe in some states you require a life-saving operation costing $\$10,000$. Then we can write utility of the form

$$\sum_{s=1}^{S} \pi_s u_s(x_s).$$

Subjective Probability

In many applications, people can agree on the set of states of nature but disagree on the probabilities. Then two DMs can consider the same state-contingent allocation to define two different lotteries.

Suppose a DM maximizes a utility function of the form $\sum_{s=1}^{S} \pi_s u_s(x_s)$. Can we determine her subjective probabilities from her behavior? In general, NO.

(Savage's Thm) Under the continuity and independence axioms, and under the additional assumption that preferences over money lotteries do not depend on the state, then preferences can be represented by an expected utility function of the form $\sum_{s=1}^{S} \pi_s u(x_s)$ and the subjective probabilities are uniquely determined.

A Violation of Subjective Expected Utility: Ellsberg Paradox

Suppose there are 2 urns. In urn 1, there are 50 red balls and 50 black balls. In urn 2, there are n red balls and $100 - n$ black balls, but n is not known.

If the DM "bets" correctly, she receives $100, and otherwise she receives nothing.

The typical DM strictly prefers a bet that a ball randomly drawn from urn 1 is red, to a bet that a ball randomly drawn from urn 2 is red (or black).

At the same time, the DM is indifferent between a bet of red or black from urn 1, and the DM is indifferent between a bet of red or black from urn 2.

These preferences exhibit *ambiguity aversion*, and are inconsistent with having subjective probabilities.