Comparison of Payoff Distributions in Terms of Return and Risk

Preliminaries

- We treat, for convenience, money as a continuous variable when dealing with monetary outcomes.
- Strictly speaking, the derivation of the expected utility representation assumed a finite number of outcomes; however, the framework can be extended, with some minor complications, to the case of an infinite domain.
- Suppose that we denote amounts of money by the continuous variable x.
- We can describe a monetary lottery by means of a *cumulative distribution function* $F : \mathbb{R} \to [0, 1]$. That is, for any x, $F(x)$ is the probability that the realized payoff is less than or equal to x.

Preliminaries (Cont.)

- Note that if the distribution function of a lottery has a density function $f(\cdot)$ associated with it, then $F(x) = \int_{-\infty}^{x} f(t)\mathrm{d}t$ for all x.
- The advantage of a formalism based on distribution functions over one based on density functions, is that the former is completely general; it does not exclude *a priori* the possibility of a discrete set of outcomes.
- From this point on, we shall work with distribution functions to describe lotteries over monetary outcomes.
- We take the lottery space $\mathcal{L}$ to be the set of all distribution functions over nonnegative amounts of money, or, more generally, over an interval $[\alpha, +\infty]$.

Preliminaries (Cont.)

- Assume a decision maker who has rational preferences $\succeq$ defined over $\mathcal{L}$. The application of the Expected Utility form to outcomes defined by a continuous variable tells us that under the assumptions of the framework, there is an assignment of utility values $u(x)$ to nonnegative amounts of money with the property that any $F(\cdot)$ can be evaluated by a utility function $U(\cdot)$ of the form:

$$U(F) = \int u(x)\mathrm{d}F(x). \tag{1}$$

- Expression (1) is the exact extension of the expected utility form to the current setting.

- Recall that we distinguish between the vNM utility function, $U(\cdot)$, defined on lotteries, and the Bernoulli utility function, $u(\cdot)$, defined on sure amounts of money.

Risk Aversion

- From now onwards, it makes sense in the current monetary context to postulate that $u(\cdot)$ is increasing and continuous.

- We now concentrate on the important property of risk aversion, its formulation in terms of the Bernoulli utility function $u(\cdot)$, and its measurement.

  DEFINITION 1. A decision maker is a *risk averter* (or exhibits *risk aversion*) if for any lottery $F(\cdot)$, the degenerate lottery that yields the amount $\int x \mathrm{d}F(x)$ with certainty is at least as good as the lottery $F(\cdot)$ itself. If the decision maker is always (i.e. for any $F(\cdot)$) indifferent between these two lotteries, we say that he is *risk neutral*.

Risk Aversion (Cont.)

- If preferences admit an expected utility representation with Bernoulli utility function $u(x)$, it follows directly from the definition of risk aversion that the decision maker is risk averse if and only if

$$\int u(x)\mathrm{d}F(x) \leq u\Big(\int x\mathrm{d}F(x)\Big) \text{ for all } F(\cdot). \qquad (2)$$

- Inequality (2) is called *Jensen's Inequality*, and it is the defining property of a concave function.

- In the context of Expected Utility framework, we see that risk aversion is equivalent to the concavity of $u(\cdot)$.

Risk Aversion (Cont.)

- Given a Bernoulli utility function, we have the following definition.

  DEFINITION 2. The *certainty equivalent* of $F(\cdot)$, denoted $c(F, u)$, is the amount of money, which makes the individual indifferent between the gamble $F(\cdot)$ and the certain amount $c(F, u)$; that is,

$$u(c(F, u)) = \int u(x) \mathrm{d}F(x). \tag{3}$$

Risk Aversion (Cont.)

PROPOSITION 1. Suppose a decision maker is an expected utility maximizer with a Bernoulli utility function $u(\cdot)$ on amounts of money. Then the following properties are equivalent:

(i) The decision maker is risk averse;

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

(iii) $c(F, u) \leq \int x \mathrm{d}F(x)$ for all $F(\cdot)$.

Measurement of Risk Aversion

DEFINITION 3. Given a (twice-differentiable) Bernoulli utility function $u(\cdot)$ for money, the *Arrow-Pratt coefficient of absolute risk aversion* at x is defined as

$$r_A(x) = \frac{-u''(x)}{u'(x)}. \tag{4}$$

- Thus, the degree of risk aversion is related to the curvature of $u(\cdot)$ (i.e. $u''(x)$).
- However, this is not an adequate measure because it is not invariant to positive linear transformations of the utility function at x.
- To make it invariant, the simplest modification is to use $\frac{u''(x)}{u'(x)}$. We also change the sign to have a positive number for an increasing and concave $u(\cdot)$.

Example

Consider the utility function $u(x) = -e^{-\alpha x}$ for $\alpha > 0$. Then $u^{'}(x) = \alpha e^{-\alpha x}$ and $u^{''}(x) = -\alpha^2 e^{-\alpha x}$. Therefore,

$$r_A(x) = \frac{-(-\alpha^2 e^{-\alpha x})}{\alpha e^{-\alpha x}} = \alpha \text{ for all } x.$$

Measurement of Risk Aversion (Cont.)

Once we are equipped with a measure of risk aversion, we can put it to use in comparative statics exercises. Two common situations are comparisons of risk attitudes *across individuals* with different utility functions and the comparisons of risk attitudes for one individual *at different levels of wealth*.

Measurement of Risk Aversion (Cont.)

Given two Bernoulli utility functions $u_1(\cdot)$ and $u_2(\cdot)$, when can we say that $u_2(\cdot)$ is unambiguously more risk averse than $u_1(\cdot)$? Several possible approaches to a definition seem plausible.

(i) $r_A(x, u_2) \geq r_A(x, u_1)$ for every x.

(ii) There exists an increasing concave function $\psi(\cdot)$ such that $u_2(x) = \psi(u_1(x))$ at all x; that is, $u_2(\cdot)$ is a concave transformation of $u_1(\cdot)$ ($u_2(\cdot)$ is "more concave" than $u_1(\cdot)$).

(iii) $c(F, u_2) \leq c(F, u_1)$ for any $F(\cdot)$.

These three definitions are equivalent.

Measurement of Risk Aversion (Cont.)

- It is common contention that wealthier people are willing to bear more risk than poorer people. Although this might be due to differences in utility functions across people, it is more likely that the source of the difference lies in the possibility that richer people can afford to take a chance.

Definition 4. The Bernoulli utility function $u(\cdot)$ for money exhibits *decreasing absolute risk aversion* if $r_A(x, u)$ is a decreasing function of x.

- Individuals whose preferences satisfy the decreasing absolute risk aversion property take more risk as they become wealthier.

Measurement of Risk Aversion (Cont.)

- The assumption of decreasing absolute risk aversion yields many other economically reasonable results concerning risk-bearing behavior.

- However, in applications, it is often too weak and, because of its analytical convenience, it is sometimes complemented by a stronger assumption: *nonincreasing relative risk aversion*.

- To understand the concept of relative risk aversion, note that the concept of absolute risk aversion is suited to the comparison of attitudes towards risky projects whose outcomes are *absolute gains or losses* from current wealth. But it is also of interest to evaluate risky projects whose outcomes are *percentage gains or losses* of current wealth.

Measurement of Risk Aversion (Cont.)

DEFINITION 5. Given a Bernoulli utility function $u(\cdot)$, the *coefficient of relative risk aversion* at x is

$$r_R(x, u) = \frac{-xu''(x)}{u'(x)}. \tag{5}$$

- Consider now how this measure varies with wealth. The property of nonincreasing relative risk aversion says that the individual becomes less risk averse with regard to gambles that are proportional to his wealth as his wealth increases.
- This is a stronger assumption than decreasing absolute risk aversion, since $r_R(x, u) = xr_A(x, u)$; a risk-averse individual with decreasing relative risk aversion will exhibit decreasing absolute risk aversion but the converse is not necessarily true.

Comparison of Payoff Distributions

- In contrast to the previous contexts where we compared utility functions, here, we will compare payoff distributions.
- There are two natural ways that random outcomes can be compared: according to the level of returns and according to the dispersion of returns.
- We will therefore attempt to give meaning to two ideas: that of a distribution $F(\cdot)$ yielding unambiguously higher returns than $G(\cdot)$ and that of $F(\cdot)$ being unambiguously less risky than $G(\cdot)$.
- These notions are known, respectively, by the technical terms of *first-order stochastic dominance* and *second-order stochastic dominance*.

First-Order Stochastic Dominance

- Note we restrict ourselves to distributions $F(\cdot)$ such that $F(0) = 0$ and $F(x) = 1$ for some x.
- We want to attach meaning to the expression: "The distribution $F(\cdot)$ yields unambiguously higher returns than the distribution $G(\cdot)$."
- At least two sensible criteria suggest themselves.
- First, we could test whether every expected utility maximizer who values more over less prefers $F(\cdot)$ to $G(\cdot)$.
- Alternatively, we could verify whether, for every amount of money x, the probability of getting **at least** x is higher under $F(\cdot)$ than under $G(\cdot)$.
- Fortunately, these two criteria lead to the same concept.

First-Order Stochastic Dominance (Cont.)

DEFINITION 6. The distribution $F(\cdot)$ first-order stochastically dominates $G(\cdot)$ if, for every nondecreasing function $u : \mathbb{R} \to \mathbb{R}$, we have

$$\int u(x)dF(x) \geq \int u(x)dG(x). \tag{6}$$

First-Order Stochastic Dominance (Cont.)

PROPOSITION 2. The distribution of monetary payoffs $F(\cdot)$ first-order stochastically dominates the distribution $G(\cdot)$ if and only if $F(x) \leq G(x)$ for every x.

- Note that first-order stochastic dominance does *not* imply that every possible return of the superior distribution is larger than every possible return of the inferior distribution (i.e. the set of possible outcomes could be the same in the two distributions).
- Although $F(\cdot)$ first-order stochastically dominating $G(\cdot)$ implies that the mean of x under $F(\cdot)$, $\int x dF(x)$, is greater than its mean under $G(\cdot)$, a ranking of the means of the two distributions does *not* imply that one first-order stochastically dominates the other; rather, the entire distribution matters.

Second-Order Stochastic Dominance

- Next, we introduce a comparison based on relative riskiness or dispersion. To avoid confusing this issue with the trade-off between returns and risk, we will restrict ourselves to comparing distributions with the same mean.

- Once again, a definition suggests itself: Given two distributions $F(\cdot)$ and $G(\cdot)$ with the same mean (i.e. $\int x dF(x) = \int x dG(x)$), we say that $G(\cdot)$ is riskier than $F(\cdot)$ if every risk averter prefers $F(\cdot)$ to $G(\cdot)$.

Definition 7. For any two distributions $F(\cdot)$ and $G(\cdot)$ with the same mean, $F(\cdot)$ second-order stochastically dominates (or is less risky than) $G(\cdot)$ if, for every nondecreasing concave function $u : \mathbb{R}_+ \to \mathbb{R}$, we have

$$\int u(x)dF(x) \geq \int u(x)dG(x). \tag{7}$$

Mean-Preserving Spreads

- Consider the following compound lottery.
- In the first stage, we have a lottery over x distributed according to $F(\cdot)$.
- In the second stage, we randomize each possible outcome x further so that final payoff is $x + z$, where z has a distribution function $H_x(z)$ with a mean of zero (i.e. $\int z dH_x(z) = 0$).
- Thus, the mean of $x + z$ is x. Let the resulting reduced lottery be denoted by $G(\cdot)$.
- When lottery $G(\cdot)$ can be obtained from lottery $F(\cdot)$ in this manner for some distribution $H_x(\cdot)$, we say that $G(\cdot)$ is a *mean-preserving spread* of $F(\cdot)$.

Example

$F(\cdot)$ is an even probability distribution between \$2 and \$3. In the second step, we may spread the \$2 outcome to an even probability between \$1 and \$3, and the \$3 outcome to an even probability between \$2 and \$4. Then $G(\cdot)$ is the distribution that assigns probability $\frac{1}{4}$ to the four outcomes: 1, 2, 3, 4. Thus, we say that $G(\cdot)$ is a mean-preserving spread of $F(\cdot)$ (and $F(\cdot)$ second-order stochastically dominates $G(\cdot)$ as indicated next).

Second-Order Stochastic Dominance (Cont.)

PROPOSITION 3. Consider two distributions $F(\cdot)$ and $G(\cdot)$ with the same mean. Then the following statements are equivalent:

(i) $F(\cdot)$ second-order stochastically dominates $G(\cdot)$.

(ii) $G(\cdot)$ is a mean-preserving spread of $F(\cdot)$.