Contents

LECTURE

Choice Under Risk

There is an early history of economic thought on risk-taking behavior, in the work of Bernoulli (1736), Fisher (1930), Keynes (1921), Menger (1934), Knight (1921), and Ramsey (1931), as well as important developments by Friedman and Savage (1948), Marschak (1950), and Arrow (1951) that parallel the von Neumann-Morganstern contribution. This concentrated on choice among lotteries, but the ideas spread to other decision-making situations. What has come to be known as Behavioral Decision Theory had its origins in the von Neumann and Morgenstern (1947) treatise on choice under uncertainty and game theory. In the following two decades, behavioral science and cognitive psychology came of age, with the participation of notable economists such as Allais (1953), Chipman (1960), Marschak (1950), Papandreau (1960), and Simon (1959).

Problematic Evidence: **?** found that 70 percent of subjects report that they would prefer a 3/4 chance of losing nothing and 1/4 chance of losing $6,000 to a 2/4 chance of losing nothing and 1/4 chance each of losing $4,000$ or $2,000$. Because the preferred lottery here is a mean-preserving spread of the less-preferred lottery, the responses of 70 percent of the subjects are inconsistent with the standard concavity assumptions.

1.1 Fair Odds

In the lifetime-consumption lottery, a consumer is endowed with an uncertain outcome for lifetime consumption. Consider a world with two states, state 1 and state 2, with probabilities p_1 and $p_2 = 1 - p_1$. The consumer's consumption outcome depends on the state of the world.

The consumer faces a consumption lottery with only two possible outcomes: outcome c_1^e with probability p_1, and outcome c_2^e with probability $1-p_1$. Equivalently, the consumer has a state-dependent consumption endowment, $\boldsymbol{c} = \langle c_1^e, c_2^e \rangle$, with an associated probability tuple $\boldsymbol{p} = \langle p_1, 1 - p_1 \rangle$. With this uncertain consumption endowment, the mean consumption outcome is therefore

$$\begin{aligned} \bar{\boldsymbol{c}}^e &\stackrel{\text{def}}{=} \boldsymbol{p} \cdot \boldsymbol{c}^e = \sum_i p_i c_i^e \\ &= p_1 c_1^e + (1 - p_1) c_2^e \end{aligned} \tag{1.1}$$

Find all the values of $\langle c_1, c_2 \rangle$ that have this same mean. These satsify the following equality.

$$\begin{aligned} \bar{\boldsymbol{c}}^e &= \boldsymbol{p} \cdot \boldsymbol{c} = \sum_i p_i c_i \\ &= p_1 c_1 + (1 - p_1) c_2 \end{aligned} \tag{1.2}$$

This is the equation of a straight line. Use this equation to solve for c_2 in terms of c_1.

$$c_2 = -\frac{p_1}{1 - p_1} c_1 + \frac{\bar{\boldsymbol{c}}^e}{1 - p_1} \tag{1.3}$$

So if we plot c_2 as a function of c_1, the slope is $-p_1/(1 - p_1)$.

The ***odds*** of an outcome is the ratio of the probability it occurs to the probability it does not occur. So the odds of outcome c_1^e is $p_1/(1 - p_1)$. The slope of this line is minus the odds of c_1^e.

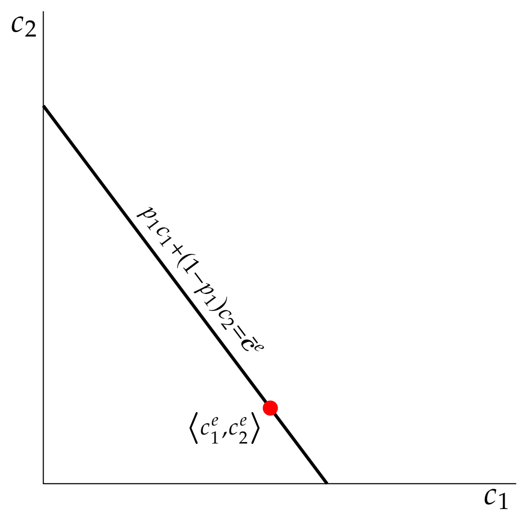

Figure 1.1: Fair-Odds Locus

Odds can be called “fair” when the expected gain from swapping lotteries is zero. In this sense, the fair-odss line represents all of the consumption lotteries that are fair relative to the endowment.

For the moment, measure how uncertain we are about the outcome by the variance. Compute the variance of any point $\langle c_1, c_2 \rangle$ on the fair-odds line.

$$\text{var}[\boldsymbol{c}] \stackrel{\text{def}}{=} \boldsymbol{p} \cdot (\boldsymbol{c} - \mu)^2 = \sum_i p_i (c_i - \mu)^2 \tag{1.4}$$

By inspection, this is minimized (at 0) by setting $\langle c_1, c_2 \rangle = \langle \mu, \mu \rangle$. This point is clearly on the fair-odds line: $p_1\mu + (1 - p_1)\mu = \mu$.

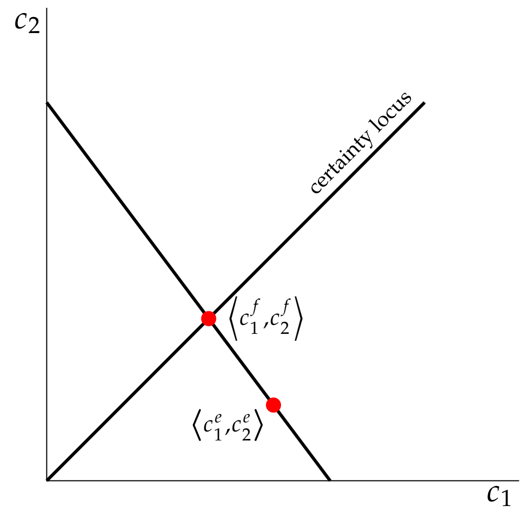

Figure 1.2: Minimum Variance
The point of intersection of the fair-odds locus and the certainty locus has the same expected return but no uncertainty.

A ***risk neutral*** consumer is indifferent between all lotteries with the same mean. However a ***risk averse*** consumer prefers less uncertainty about outcomes. Risk aversion means that among all bundles with equal expected value, you prefer the sure thing. It follows immediately that *if* the fair-odds locus also represents the available market trade-offs, risk averse individuals will eliminate all risk by trading in the market along this tradeoff until they reach the certainty locus. So a risk-averse consumer can use the certainty locus to find the most preferred point on any fair odds locus.

Imagine we can draw indifference curves for this consumer in $\langle c_1, c_2 \rangle$ space. The curve that passes through $\langle \bar{c}, \bar{c} \rangle$ must stay above the fair odds locus at all other points. So risk aversion has a simple graphical corollary.

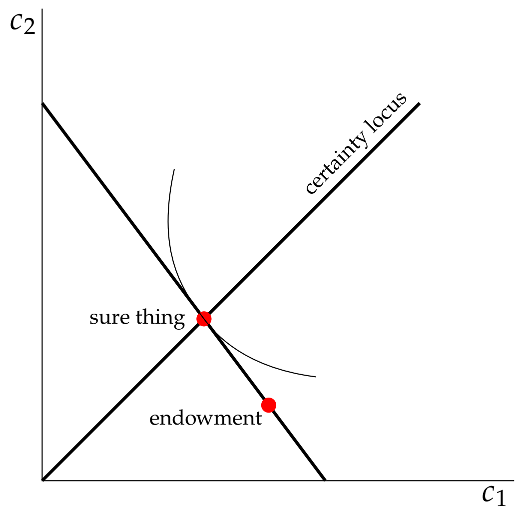

Figure 1.3: Risk Aversion
For a risk averse consumer, the indifference curve through the sure thing must lie above the rest of the fair-odds locus.

Exercise 1.1 A risk-averse person with a dollar offered a game with fair odds: based on a coin flip, the dollar is lost or doubled. Explain why this game will be rejected.

1.1.1 Risky Utility

Economists think of consumers as deriving *utility* from their consumption. To make this concrete, introduce a utility function u that is continuous, strictly increasing, and strictly concave. Note how the concavity implies that expected utility $\bar{u}$ is below the utility of the mean outcome $\bar{c}^e$. That is, the concavity of the utility function implies that the consumer is risk averse.

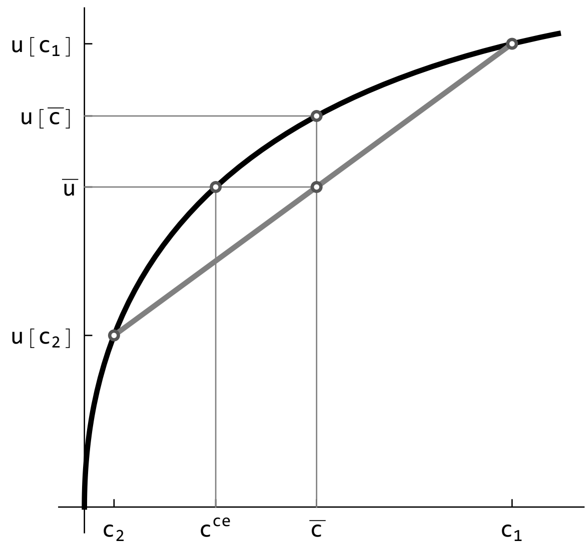

Figure 1.4: Geometry of Risk Aversion
The consumer gets c_1 with probability p_1 and c_2 with probability p_2, with a mean outcome of $\bar{c}$. The certainty equivalent is c^{ce}, and $\bar{c} - c^{ce}$ is the maximum risk premium the consumer could pay just to elminate all risk.

Remark 1.1 That u is continuous implies the certainty equivalent exists (by the intermediate value theorem). That u is also strictly increasing ensures that the certainty equivalent is unique. That u is also strictly concave ensures that the certainty equivalent is less than mean consumption.

Exercise 1.2 Consider a utility function $u : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ that is differentiable, strictly increasing, and strictly concave. Prove that if u assigns one lottery ($\boldsymbol{p}$) higher utility than another($\boldsymbol{q}$), then for $b > 0$, so does the affine tranformation $a + bu$. Can you find a strictly increasing transformation $t : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ that will not do this?

1.2 Insurance

If you are willing to pay something to reduce your risk, you may look for insurance. An insurance contract has a premium that you must pay in advance, but in return you can get payout in the bad state of the world.

Suppose c_1 is the real value of consumption in the good state of the world, but in the bad state of the world it is $c_2 = c_1 - \Delta$. Here $\Delta = c_1 - c_2$ is a possible loss to the consumer. Now introduce the cost (δ) of fully insuring against this loss (i.e., the cost of certainty). The consumer may choose to partially insure (or even to over insure). If she pays an insurance premium $\alpha\delta$ and incurs loss Δ, she receives the indemnity $\alpha\Delta$ from the insurance company.

Table 1.1

Probability	Outcome
p_1	$c_1 - \alpha\delta$
$1 - p_1$	$c_2 - \alpha\delta + \alpha\Delta$

Comment: The contract is actuarially fair if $\delta = (1 - p_1)\Delta$, so that the premium equals the expected value of the payout.

Assume the consumer maximizes *expected* utility, and consider the following *insurance problem*.

$$\max_{\alpha} p_1\, u[c_1 - \alpha\delta] + (1 - p_1)\, u[c_2 - \alpha\delta + \alpha\Delta] \tag{1.5}$$

To produce the first order necessary condition (FONC) for an extremum, set the first-order derivative to 0. You should use your standard rules of differentiation, including the sum rule and the chain rule.

$$-\delta p_1\, u'[c_1 - \alpha\delta] + (1 - p_1)(\Delta - \delta)\, u'[c_2 - \alpha\delta + \alpha\Delta] = 0 \tag{1.6}$$

Equivalently, since $c_2 = c_1 - \Delta$,

$$-\delta p_1\, u'[c_1 - \alpha\delta] + (1 - p_1)(\Delta - \delta)\, u'[c_1 - \alpha\delta - (1 - \alpha)\Delta] = 0 \tag{1.7}$$

In this expression, the derivative function u' is evaluated at two different places: the state 1 outcome $c_1 - \alpha\delta$, and that state 2 outcome $c_1 - \alpha\delta - (1 - \alpha)\Delta$.

Note: If the contract is actuarially fair, then $p_1\delta = (1 - p_1)(\Delta - \delta)$ so the FONC becomes

$$u'[c_1 - \alpha\delta] = u'[c_1 - \alpha\delta - (1 - \alpha)\Delta] \tag{1.8}$$

This requires two equal values of u'. Since u is strictly concave, u' is strictly decreasing, so the FONC requires the two inputs to be equal. So the consumer must set $\alpha = 1$ to satisfy the FONC.

When insurance is actuarially fair, the result is complete income smoothing!