Part III: Uncertainty and Strategy

7. Uncertainty
8. Game Theory

Chapter 7
Uncertainty
Part I

Ming-Ching Luoh

2020.11.27.

Mathematical Statistics

- *Random variable:* A variable that records, in numerical form, the possible outcomes from some random event.
- *Probability density function (PDF):* A function $f(x)$ that shows the probabilities associated with the possible outcomes from a random variable.
- *Expected value of a random variable:* The outcome of a random variable that will occur "on average." The expected value is denoted by $E(x)$.

  If x is a discrete random variable with n outcomes, then

$$E(x) = \sum_{i=1}^{n} x_i f(x_i).$$

- If x is a continuous random variable, then

$$E(x) = \int_{-\infty}^{+\infty} x f(x) dx.$$

- *Variance and standard deviation of a random variable:* These concepts measure the dispersion of a random variable about its expected value.

  In the discrete case,

$$Var(x) = \sigma_x^2 = \sum_{i=1}^{n} [x_i - E(x)]^2 f(x_i).$$

  In the continuous case,

$$Var(x) = \sigma_x^2 = \int_{-\infty}^{+\infty} [x - E(x)]^2 f(x) dx.$$

  The standard deviation is the square root of the variance.

Fair Gambles and the Expected Utility Hypothesis

- A “fair" gamble is a specified set of prizes and associated probabilities that has an expected value of zero.
- It has long been recognized that people would prefer not to play fair games. For example, people tend to refuse the gamble of winning $1 million with probability 1/2 and losing $1 million with probability 1/2.
- Daniel Bernoulli’s famous study of St. Petersburg paradox in 18th century provided the starting point for virtually all studies of the behavior of individuals in uncertain situations.

St. Petersburg paradox

In the St. Petersburg paradox, the following gamble is proposed.

- A coin is flipped until a head appears.
- If a head appears on the nth flip, the player is paid $\$2^n$. If x_i represents the prize awarded when the first head appears on the ith trial, then

$$x_1 = \$2, x_2 = \$4, x_3 = \$8, \cdots, x_n = \$2^n.$$

- The probability of getting a head on the ith trial is $\left(\frac{1}{2}\right)^i$, hence the probability of the prizes given in the ith trial is

$$\pi_1 = \frac{1}{2}, \pi_2 = \frac{1}{4}, \pi_3 = \frac{1}{8}, \cdots, \pi_n = \frac{1}{2^n}.$$

- Therefore, the expected value of the gamble is infinite:

$$E(x) = \sum_{i=1}^{\infty} \pi_i \cdot x_i = \sum_{i=1}^{\infty} \frac{1}{2^i} \cdot 2^i = 1 + 1 + 1 + \cdots + 1 + \cdots = \infty.$$

- Because no player would pay a lot to play this game, this is then the paradox: Bernoulli's gamble is in some sense not worth its (infinite) expected dollar value.

- This paradox is named after the city where Bernoulli's original manuscript was published. The article has been reprinted as D. Bernoulli, "Exposition of a New Theory on the Measurement of Risk," *Econometrica* 22 (January 1954): 23-36.

Expected Utility

- Bernoulli's solution to this paradox was to argue that individuals do not care directly about the dollar values of the prizes. They care about the utility that the dollars provide.
- If we assume diminishing marginal utility of wealth, the St. Petersburg game may converge to a finite expected utility value even though its expected monetary value is infinite.
- Because the gamble only provides a finite expected utility, individuals would only be willing to pat a finite amount to play it.

Example 7.1 Bernoulli's Solution to the Paradox and Its Shortcomings

- Suppose that the utility of each prize is given by

  $$U(x_i) = \ln x_i$$

  This utility function exhibits diminishing marginal utility (i.e. $U' > 0$ but $U'' < 0$), and the expected utility value converges to a finite number:

  $$\begin{aligned} \text{expected utility} &= \sum_{i=1}^{\infty} \pi_i U(x_i) = \sum_{i=1}^{\infty} \frac{1}{2^i} \ln(2^i) \\ &= \ln 2 \sum_{i=1}^{\infty} \frac{i}{2^i} = 2 \ln 2 \approx 1.39 \end{aligned}$$

- Thus, assuming that the large prizes promised by the St. Petersburg paradox encounter diminishing marginal utility permitted Bernoulli to offer a solution to the paradox.

- Unfortunately, Bernoulli's solution to the St. Petersburg paradox does not completely solve the problem.
- As long as there is no upper bound to the utility function, the paradox can be regenerated by redefining the gamble's prizes.
- For example, prizes can be set as $x_i = e^{2^i}$, in which case

$$U(x_i) = \ln e^{2^i} = 2^i$$

and the expected utility from the gamble would be infinite.
- The prizes in this redefined gamble are large. For example, if a head first appears on the 5th flip, a person would win $e^{2^5} =$ \$79 trillion, although the probability of winning would be only $\frac{1}{2^5} = 0.031$.
- This gamble still seems to be unlikely. Hence the St. Petersburg game remains a paradox.

The von Neumann-Morgenstern Theorem

- In their book *The Theory of Games and Economic Behavior*, John von Neumann and Oscar Morgenstern developed a mathematical foundation for Bernoulli's solution to the St. Petersburg paradox.
- They laid out basic axioms of rationality and showed that any person who is rational in this would make choices under uncertainty as though he or she had a utility function over money $U(x)$ and maximized the expected value of $U(x)$, rather than the expected value of x itself.

The von Neumann-Morgenstern utility index

- Let the prizes be denoted by $x_1, x_2, \cdots x_n$, and assume that these have been arranged in order of ascending desirability.
- Assign arbitrary utility function numbers to these two extreme prizes such as

$$U(x_1) = 0,$$
$$U(x_n) = 1.$$

- The point of the von Neumann-Morgenstern is to show that a reasonable way exists to assign specific utility numbers to the other prizes available.

- Consider any other prize x_i. Ask the individual to state the probability π_i at which he or she would be indifferent between x_i with *certainty*, and a *gamble* offering prizes of x_n with probability π_i and x_1 with probability $1 - \pi$.
- It seems reasonable that such a probability will exist.
- The probability π_i measures how desirable the prize x_i is.
- The von Neumann-Morgenstern technique defines the utility of x_i as the expected utility of the gamble that the individual consider equally desirable to x_i:

$$\begin{aligned} U(x_i) &= \pi_i U(x_n) + (1 - \pi_i) U(x_1) \\ &= \pi_i \cdot 1 + (1 - \pi_1) \cdot 0 = \pi_i \end{aligned}$$

- The utility index attached to any other prize is simply the probability of winning the top prize in a gamble the individual regards as equivalent to the prize in question.

Expected utility maximization

- Suppose that a utility index π_i has been assigned to every prize x_i, with $\pi_1 = 0, \pi_n = 1$.
- Using these utility indices, we can show that a "rational" individual will choose among gambles based on their expected "utilities".
- Consider two gambles. Gamble A offers x_2 with probability a and x_3 with probability $1 - a$. Gamble B offers x_4 with probability b and x_5 with probability $1 - b$.

$$\begin{aligned} \text{expected utility of A} &= E_A[U(x)] = aU(x_2) + (1-a)U(x_3), \\ \text{expected utility of B} &= E_B[U(x)] = bU(x_4) + (1-b)U(x_5). \end{aligned}$$

- Substituting the utility index numbers gives

$$\begin{aligned} E_A[U(x)] &= a\pi_2 + (1-a)\pi_3, \\ E_B[U(x)] &= b\pi_4 + (1-b)\pi_5. \end{aligned}$$

- We want to show that the individual will prefer gamble A to gamble B if and only if

$$E_A[U(x)] > E_B[U(x)]$$

- Since the individual is indifferent between x_2 and a gamble promising x_1 with probability $1-\pi_2$ and x_n with probability π_2, the expected utility of gamble A is

$$\begin{aligned} E_A[U(x)] &= a\pi_2 + (1-a)\pi_3 = a[(1-\pi_2)U(x_1) + \pi_2 U(x_n)] \\ &+ (1-a)[(1-\pi_3)U(x_1) + \pi_3 U(x_n)] \\ &= [a\pi_2 + (1-a)\pi_3]U(x_n) \\ &+ [a(1-\pi_2) + (1-a)(1-\pi_3)]U(x_1) \end{aligned}$$

- Therefore, gamble A is equivalent to a gamble promising x_n with probability $a\pi_2 + (1 - a)\pi_3$, and gamble B is equivalent to a gamble promising x_n with probability $b\pi_4 + (1 - b)\pi_5$.
- The individual will choose gamble A if and only if

$$a\pi_2 + (1 - a)\pi_3 > b\pi_4 + (1 - b)\pi_5.$$

  This is exactly the condition that $E_A[U(x)] > E_B[U(x)]$.
- An individual will choose the gamble that provides the highest level of expected utility.
- **Expected utility maximization.** If individuals obey the von Neumann-Morgenstern axioms of behavior in uncertain situations, they will act as if they choose the option that maximizes the expected value of their von Neumann-Morgenstern utility.

Risk Aversion

- Economists have found that people tend to avoid risky situations, even if the situation amount to a fair gamble.
- Extra money may provide people with decreasing marginal utility.
- Starting from a wealth of $50,000, the individual would be reluctant to take a $10,000 bet on a coin flip because the 50% chance of the increased utility does not compensate for the 50% chance of decreased utility.
- On the other hand, a bet of only $1 on a coin flip is relatively inconsequential.

Risk aversion and fair gambles

Figure 7.1 Utility of Wealth Facing a Fair Bet

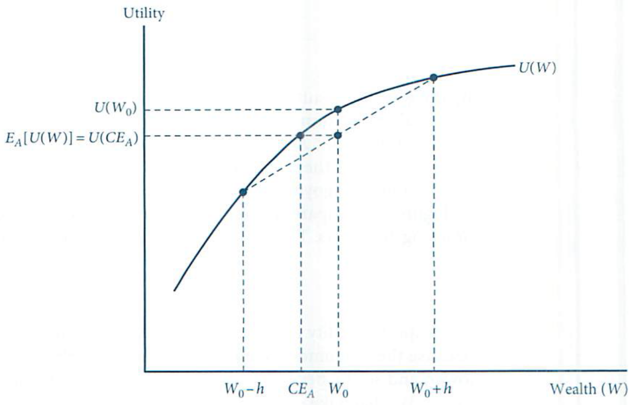

- W_0 represents an individual's current wealth and $U(W)$ is a von Neumann-Morgenstern utility function.

- $U(W)$ is drawn as a concave function to reflect the assumption of diminishing marginal utility of wealth.

- The expected utility of participating a fair gamble A, which is a 50-50 chance of winning or losing h dollars, is

$$E_A[U(W)] = \frac{1}{2}U(W_0 + h) + \frac{1}{2}U(W_0 - h).$$

- It is clear from the geometry of the figure that

$$U(W_0) > E_A[U(W)].$$

  This person will prefer to keep his or her current wealth rather than taking the fair gamble because winning a fair bet adds to enjoyment less than losing hurts.

Figure 7.2 Comparing Two Fair Bets of Differing Variability

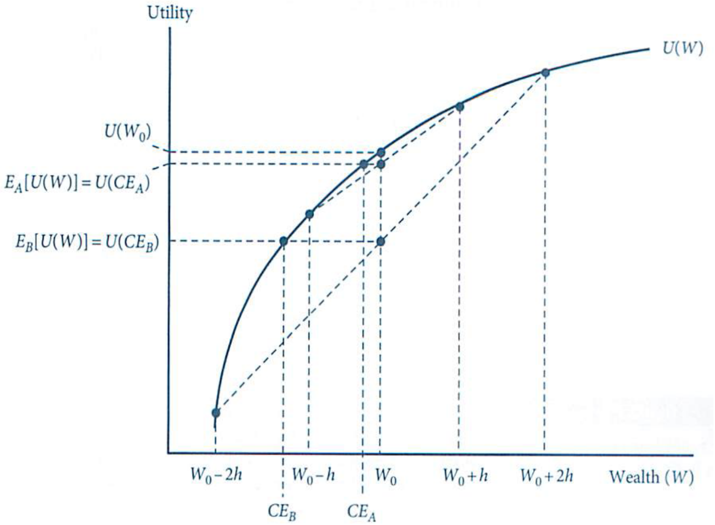

- Figure 7.2 compares gamble A to a new gamble, B, which is a 50-50 chance of winning or losing $2h$ dollars. Expected utility from gamble B equals

$$E_B[U(W)] = \frac{1}{2}U(W_0 + 2h) + \frac{1}{2}U(W_0 - 2h)$$

- Because the outcomes are more variable in gamble B than A, the expected utility of B is lower, and so the person prefers A to B (although he or she would prefer to keep initial wealth W_0 than take either gamble).

Risk aversion and insurance

- Note that in Figure 7.2, a certain wealth of CE_A provides the same expected utility as does participating in gamble A. CE_A is referred to as the *certainty equivalent* of gamble A.
- The individual would be willing to pay up to $W_0 - CE_A$ to avoid participating in the gamble. This explains why people buy insurance.
- The person in Figure 7.2 would pay even more to avoid taking the larger gamble, B, as shown by the observation that $W_0 - CE_B > W_0 - CE_A$ in the figure.
- **Risk aversion.** An individual who always refuses fair bets is said to be risk averse. If individuals exhibit a diminishing marginal utility of wealth, they will be risk averse. As a consequence, they will be willing to pay something to avoid taking fair bets.

Example 7.2 Willingness to Pay for Insurance

- Consider a person with a current wealth of \$100,000 who faces a 25% chance of losing his automobile worth \$20,000 through theft during the next year.
- Suppose that this person's Von Neumann-Morgenstern utility function is, $U(W) = \ln(W)$.
- This person's expected utility without insurance will be

$$\begin{aligned} E_{no}[U(W)] &= 0.75U(100,000) + 0.25U(80,000) \\ &= 0.75 \ln 100,000 + 0.25 \ln 80,000 = 11.45714. \end{aligned}$$

- In this situation, a fair insurance premium would be \$5,000 $(25\% \times 20,000)$. The expected utility of fair insurance is

$$E_{fair}[U(W)] = U(95,000) = \ln 95,000 = 11.46163.$$

- This person is made better off by purchasing fair insurance.
- The maximum insurance premium (x) he or she would be willing to pay can be determined by solving the following equation.

$$\begin{aligned} E_{wtp}[U(W)] &= U(100,000 - x) \\ &= \ln(100,000 - x) = 11.45714 \end{aligned}$$

Therefore,

$$\begin{aligned} 100,000 - x &= e^{11.45714}, \\ x &= 5,426. \end{aligned}$$

Measuring Risk Aversion

- The most commonly used measure of risk aversion was developed by J. W. Pratt in the 1960s. It is defined as

$$r(W) = -\frac{U''(W)}{U'(W)}$$

  Because $U''(W) < 0$ from a diminishing marginal utility of wealth, $r(W)$ is positive.

- This measure is not affected by which particular von Neumann-Morgenstern ordering is used.

Risk aversion and insurance premiums

- Suppose the winnings from a fair bet are denoted by the random variable h, with $E(h) = 0$.
- Let p be the size of the insurance premium that would make the individual indifferent between taking the fair bet h and paying p with certainty to avoid the gamble:

$$E[U(W + h)] = U(W - p),$$

where W is the individual's current wealth.

- Expand both sides of the equation using Taylor's series.
- Because p is a fixed amount, a linear approximation to the right side of the equation will suffice:

$$U(W - p) = U(W) - pU'(W) + \text{ higher} - \text{order terms}$$

- For the left side, we need a quadratic approximation to allow for the variability in the gamble, h:

$$\begin{aligned} E[U(W+h)] &= E\left[U(W) + hU'(W) + \frac{h^2}{2}U''(W) + \text{ higher} - \text{order terms}\right] \\ &= U(W) + E(h)U'(W) + \frac{E(h^2)}{2}U''(W) + \text{ higher} - \text{order terms.} \end{aligned}$$

- Recall $E(h) = 0$, let $k = \frac{E(h^2)}{2}$ and drop the higher-order terms, we have

$$\begin{aligned} U(W) - pU'(W) &\approx U(W) + kU''(W)] \\ p &\approx -\frac{kU''(W)}{U'(W)} = kr(W) \end{aligned}$$

- The amount that a risk-averse individual is willing to pay to avoid a fair bet is approximately proportional to Pratt's risk aversion measure.

- An important question is whether risk aversion increase or decreases with wealth. It depends on the precise shape of the utility function.

- If utility is quadratic in wealth,

$$U(W) = a + bW + cW^2,$$

where $b > 0$ and $c < 0$, then Pratt's risk aversion measure is

$$r(W) = -\frac{U''(W)}{U'(W)} = \frac{-2c}{b + 2cW},$$

which increases as wealth increases because

$$\frac{\partial r(W)}{\partial W} = \frac{4c^2}{(b + 2cW)^2} > 0$$

- If utility is logarithmic in wealth, $U(W) = \ln W$, then we have

$$r(W) = -\frac{U''(W)}{U'(W)} = \frac{1}{W},$$

which decreases as wealth increases.

- The exponential utility function

$$U(W) = -e^{-AW}$$

(where A is a positive constant) exhibit constant absolute risk aversion over all rangees of wealth because

$$r(W) = -\frac{U''(W)}{U'(W)} = \frac{A^2 e^{-AW}}{Ae^{-AW}} = A.$$

This feature of the exponential utility function can be used to provide numerical estimate of the willingness to pay to avoid gambles.

Example 7.3 Constant Risk Aversion

- Suppose an individual whose initial wealth is W_0 and whose utility function exhibits constant absolute risk aversion is facing a fair gamble of \$1,000. How much (f) would he or she pay to avoid the risk?
- To find f, we set up the following equation:

$$
\begin{aligned}
-e^{-A(W_0-f)} &= -\frac{1}{2}e^{-A(W_0+1{,}000)} - \frac{1}{2}e^{-A(W_0-1{,}000)}, \\
\text{or } e^{Af} &= \frac{1}{2}e^{-1{,}000A} + \frac{1}{2}e^{1{,}000A}
\end{aligned}
$$

- The willingness to pay to avoid a given gamble is independent of initial wealth (W_0).
- If $A = 0.0001$, then $f = 49.9$; If $A = 0.0003$, then $f = 147.8$.
- Values of the risk aversion parameter A in these ranges are sometimes used for empirical investigations.

- It wealth W is a Normal random variable with mean μ and variance σ^2. Its probability density function is

$$f(W) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-(W-\mu)^2/2\sigma^2}.$$

- If this person's utility for wealth is $U(W) = -e^{-AW}$, then the expected utility over risky wealth is

$$\begin{aligned} E[U(W)] &= \int_{-\infty}^{\infty} U(W) f(W) dW \\ &= \frac{1}{\sqrt{2\pi\sigma^2}} \int_{-\infty}^{\infty} -e^{-AW} e^{-(W-\mu)^2/2\sigma^2} dW \end{aligned}$$

- Let $z = (W - \mu)/\sigma$, then $W = \mu + \sigma z$, $dW = \sigma dz$, then

$$
\begin{aligned}
E[U(W)] &= \frac{1}{\sqrt{2\pi\sigma^2}} \int_{-\infty}^{\infty} e^{-A(\sigma z+\mu)} e^{-z^2/2} \sigma dz \\
&= \frac{1}{\sqrt{2\pi\sigma^2}} \int_{-\infty}^{\infty} e^{-(z+A\sigma)^2/2} e^{(-A\mu + A^2\sigma^2/2)} \sigma dz \\
&= e^{-A(\mu - A\sigma^2/2)} \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} e^{-(z+A\sigma)^2/2} dz \\
&= e^{-A(\mu - A\sigma^2/2)} = e^{-A} \exp(\mu - A\sigma^2/2)
\end{aligned}
$$

- This is simply the monotonic transformation of $\mu - A\sigma^2/2$. This person's preferences can be represented by

$$\mu - \frac{A}{2}\sigma^2 = CE$$

This person would be indifferent between his or her risky wealth (Normally distributed with mean μ and variance σ^2) and certain wealth with mean CE and no variance.

Relative risk aversion

- It is unlikely that the willingness to pay to avoid a given gamble is independent of a person's wealth.
- A more appealing assumption may be that such willingness to pay is inversely proportional to wealth and that the expression

$$rr(W) = Wr(W) = -W\frac{U''(W)}{U'(W)}$$

might be approximately constant. The $rr(W)$ function is a measure of *relative risk aversion*.

- The power utility function

$$U(W,R) = \begin{cases} W^R/R & \text{if } R < 1, R \neq 0 \\ \ln W & \text{if } R = 0 \end{cases}$$

exhibits diminishing absolute risk aversion,

$$r(W) = -\frac{U''(W)}{U'(W)} = -\frac{(R-1)W^{R-2}}{W^{R-1}} = \frac{1-R}{W},$$

but constant relative risk aversion ($CRRA$),

$$rr(W) = Wr(W) = 1 - R$$

- Empirical evidence is generally consistent with values of R in the range of -3 to -1.

Example 7.4 Constant Relative Risk Aversion

- An individual with a constant relative risk aversion utility function will be concerned about proportional gains or loss of wealth.
- What fraction of initial wealth (f) such a person would be willing to give up to avoid a fair gamble of, say, 10% of initial wealth.
- Assume $R = 0$, so that $U(W, R) = \ln W$. That is

$$
\begin{aligned}
\ln[(1-f)W_0] &= = 0.5\ln(1.1W_0) + 0.5\ln(0.9W_0). \\
\ln(1-f) &= 0.5[\ln 1.1 + \ln 0.9] = \ln 0.99^{0.5} \\
1-f &= 0.99^{0.5} = 0.995 \\
f &= 0.005.
\end{aligned}
$$

A person will sacrifice up to 0.5% of wealth to avoid a 10 percent gamble.

- For the case of $R = -2$, $U(W) = \frac{W^{-2}}{-2}$. Therefore

$$
\begin{aligned}
\frac{[(1-f)W_0]^{-2}}{-2} &= 0.5\frac{[1.1W_0]^{-2}}{-2} + 0.5\frac{[0.9W_0]^{-2}}{-2} \\
\frac{1}{(1-f)^2} &= \frac{0.5}{1.1^2} + \frac{0.5}{0.9^2} \\
f &= 0.015 = 1.5\%
\end{aligned}
$$

- The more risk-averse ($R = -2$ v.s. $R = 0$) person would be willing to give up 1.5% of the initial wealth to avoid a 10% gamble.