2

Bayesian Decision Theory

> In a formal model the conclusions are derived from definitions and assumptions. ... But with informal, verbal reasoning ... one can argue until one is blue in the face ... because there is no criterion for deciding the soundness of an informal argument.
>
> Robert Aumann

2.1 The Rational Actor Model

In this section we develop a set of behavioral properties, among which consistency is the most prominent, that together ensure that we can model the individual as maximizing a preference function over outcomes, subject to constraints.

A *binary relation* $\odot_{\mathrm{A}}$ on a set A is a subset of $A \times A$. We usually write the proposition $(x, y) \in \odot_{\mathrm{A}}$ as $x \odot_{\mathrm{A}} y$. For instance, the arithmetical operator "less than" ($<$) is a binary relation, where $(x, y) \in <$ is normally written $x < y$.[1] A *preference ordering* $\succeq_{\mathrm{A}}$ on A is a binary relation with the following three properties, which must hold for all $x, y, z \in A$ and any set B:

1. **Complete**: $x \succeq_{\mathrm{A}} y$ or $y \succeq_{\mathrm{A}} x$,
2. **Transitive**: $x \succeq_{\mathrm{A}} y$ and $y \succeq_{\mathrm{A}} z$ imply $x \succeq_{\mathrm{A}} z$,
3. **Independence from Irrelevant Alternatives**: For $x, y \in B$, $x \succeq_{\mathrm{B}} y$ if and only if $x \succeq_{\mathrm{A}} y$.

Because of the third property, we need not specify the choice set and can simply write $x \succeq y$. We also make the behavioral assumption that given any choice set A, the individual chooses an element $x \in A$ such that for all $y \in A$, $x \succeq y$. When $x \succeq y$, we say "x is weakly preferred to y."

[1]Additional binary relations over the set **R** of real numbers include "$>$,""$<$,""$\leq$," "$=$," "$\geq$," and "$\neq$," but "$+$" is not a binary relation, because $x + y$ is not a proposition.

Completeness implies that any member of A is weakly preferred to itself (for any x in A, $x \succeq x$). In general, we say a binary relation $\odot$ is *reflexive* if, for all x, $x \odot x$. Thus, completeness implies reflexivity. We refer to $\succeq$ as "weak preference" in contrast to $\succ$ as "strong preference." We define as $x \succ y$ to mean "it is false that $y \succeq x$." We say x and y are *equivalent* if $x \succeq y$ and $y \succeq x$, and we write $x \simeq y$. As an exercise, you may use elementary logic to prove that if $\succeq$ satisfies the completeness condition, then $\succ$ satisfies the following *exclusion* condition: if $x \succ y$, then it is false that $y \succ x$.

The second condition is *transitivity*, which says that $x \succeq y$ and $y \succeq z$ imply $x \succeq z$. It is hard to see how this condition could fail for anything we might like to call a "preference ordering."[2] As an exercise, you may show that $x \succ y$ and $y \succeq z$ imply $x \succ z$, and $x \succeq y$ and $y \succ z$ imply $x \succ z$. Similarly, you may use elementary logic to prove that if $\succeq$ satisfies the completeness condition, then $\simeq$ is transitive (that is, satisfies the transitivity condition).

When these three conditions are satisfied, we say the preference relation $\succeq$ is *consistent*. If $\succeq$ is a consistent preference relation, then there always exists a preference function such that the individual behaves as if maximizing this preference function over the set A from which he or she is constrained to choose. Formally, we say that a preference function $u : A \rightarrow \mathbf{R}$ *represents* a binary relation $\succeq$ if, for all $x, y \in A$, $u(x) \geq u(y)$ if and only if $x \succeq y$. We have:

THEOREM 2.1 *A binary relation $\succeq$ on the finite set A of payoffs can be represented by a preference function $u : A \rightarrow \mathbf{R}$ if and only if $\succeq$ is consistent.*

It is clear that $u(\cdot)$ is not unique, and indeed, we have the following.

THEOREM 2.2 *If $u(\cdot)$ represents the preference relation $\succeq$ and $f(\cdot)$ is a strictly increasing function, then $v(\cdot) = f(u(\cdot))$ also represents $\succeq$. Conversely, if both $u(\cdot)$ and $v(\cdot)$ represent $\succeq$, then there is an increasing function $f(\cdot)$ such that $v(\cdot) = f(u(\cdot))$.*

The first half of the theorem is true because if f is strictly increasing, then $u(x) > u(y)$ implies $v(x) = f(u(x)) > f(u(y)) = v(y)$, and conversely.

[2]The only plausible model of intransitivity with some empirical support is *regret theory* (Loomes 1988; Sugden 1993). This analysis applies, however, to only a narrow range of choice situations.

For the second half, suppose $u(\cdot)$ and $v(\cdot)$ both represent $\succeq$, and for any $y \in \mathbf{R}$ such that $v(x) = y$ for some $x \in X$, let $f(y) = u(v^{-1}(y))$, which is possible because v is an increasing function. Then $f(\cdot)$ is increasing (because it is the composition of two increasing functions) and $f(v(x)) = u(v^{-1}(v(x))) = u(x)$, which proves the theorem.

2.2 Time Consistency and Exponential Discounting

The central theorem on choice over time is that time consistency results from assuming that *utility be additive across time periods and the instantaneous utility function be the same in all time periods, with future utilities discounted to the present at a fixed rate* (Strotz 1955). This is called *exponential discounting* and is widely assumed in economic models. For instance, suppose an individual can choose between two consumption streams $x = x_0, x_1, \ldots$ or $y = y_0, y_1, \ldots$ According to exponential discounting, he has a utility function $u(x)$ and a constant $\delta \in (0, 1)$ such that the total utility of stream x is given by[3]

$$U(x_0, x_1, \ldots) = \sum_{k=0}^{\infty} \delta^k u(x_k). \tag{2.1}$$

We call δ the individual's *discount factor*. Often we write $\delta = e^{-r}$ where we interpret $r > 0$ as the individual's one-period, continuous-compounding *interest rate*, in which case (2.1) becomes

$$U(x_0, x_1, \ldots) = \sum_{k=0}^{\infty} e^{-rk} u(x_k). \tag{2.2}$$

This form clarifies why we call this "exponential" discounting. The individual strictly prefers consumption stream x over stream y if and only if $U(x) > U(y)$. In the simple compounding case, where the interest accrues at the end of the period, we write $\delta = 1/(1 + r)$, and (2.2) becomes

$$U(x_0, x_1, \ldots) = \sum_{k=0}^{\infty} \frac{u(x_k)}{(1 + r)^k}. \tag{2.3}$$

[3]Throughout this text, we write $x \in (a, b)$ for $a < x < b$, $x \in [a, b)$ for $a \le x < b$, $x \in (a, b]$ for $a < x \le b$, and $x \in [a, b]$ for $a \le x \le b$.

The derivation of (2.2) is a bit tedious, and except for the exponential discounting part, is intuitively obvious. So let us assume utility $u(x)$ is additive and has the same shape across time, and show that exponential discounting must hold. I will construct a very simple case that is easily generalized. Suppose the individual has an amount of money M that he can either invest or consume in periods $t = 0, 1, 2$. Suppose the interest rate is r, and interest accrues continually, so \$1 put in the bank at time $k = 0$ yields e^{rk} at time k. Thus, by putting an amount $x_k e^{-rk}$ in the bank today, the individual will be able to consume x_k in period k. By the additivity and constancy across periods of utility, the individual will maximize some objective function

$$V(x_0, x_1, x_2) = u(x_0) + au(x_1) + bu(x_2), \tag{2.4}$$

subject to the income constraint

$$x_0 + e^{-r}x_1 + e^{-2r}x_2 = M.$$

where r is the interest rate. We must show that $b = a^2$ if and only if the individual is time consistent. We form the Lagrangian

$$\mathcal{L} = V(x_0, x_1, x_2) + \lambda(x_0 + e^{-r}x_1 + e^{-2r}x_2 - M),$$

where λ is the Lagrangian multiplier. The first-order conditions for a maximum are then given by $\partial\mathcal{L}/\partial x_i = 0$ for $i = 0, 1, 2$. Solving these equations, we find

$$\frac{u'(x_1)}{u'(x_2)} = \frac{be^r}{a}. \tag{2.5}$$

Now, time consistency means that after consuming x_0 in the first period, the individuals will still want to consume x_1 in the second period and x_2 in the third. But now his objective function is

$$V(x_1, x_2) = u(x_1) + au(x_2), \tag{2.6}$$

subject to the (same) income constraint

$$x_1 + e^{-r}x_2 = (M - x_0)e^{-r},$$

We form the Lagrangian

$$\mathcal{L}_1 = V(x_1, x_2) + \lambda(x_1 + e^{-r}x_2 - (M - x_0)e^{-r}),$$

where λ is the Lagrangian multiplier. The first-order conditions for a maximum are then given by $\partial\mathcal{L}_1/\partial x_i = 0$ for $i = 1, 2$. Solving these equations, we find

$$\frac{u'(x_1)}{u'(x_2)} = ae^r. \tag{2.7}$$

Now, time consistency means that (2.5) and (2.7) are equal, which means $a^2 = b$, as required.

2.3 The Expected Utility Principle

What about decisions in which a stochastic event determines the payoffs to the players? Let X be a set of "prizes." A *lottery* with payoffs in X is a function $p{:}X\rightarrow[0, 1]$ such that $\sum_{x\in X} p(x) = 1$. We interpret $p(x)$ as the probability that the payoff is $x \in X$. If $X = \{x_1, \ldots, x_n\}$ for some finite number n, we write $p(x_i) = p_i$.

The *expected value* of a lottery is the sum of the payoffs, where each payoff is weighted by the probability that the payoff will occur. If the lottery l has payoffs $x_1 \ldots x_n$ with probabilities $p_1, \ldots, p_n$, then the expected value $\mathbf{E}[l]$ of the lottery l is given by

$$\mathbf{E}[l] = \sum_{i=1}^{n} p_i x_i.$$

The expected value is important because of the law of large numbers (Feller 1950), which states that as the number of times a lottery is played goes to infinity, the average payoff converges to the expected value of the lottery with probability 1.

Consider the lottery l_1 in pane (a) of figure 2.1, where p is the probability of winning amount a and $1-p$ is the probability of winning amount b. The expected value of the lottery is then $\mathbf{E}[l_1] = pa + (1-p)b$. Note that we model a lottery a lot like an extensive form game, except there is only one player.

Consider the lottery l_2 with the three payoffs shown in pane (b) of figure 2.1. Here p is the probability of winning amount a, q is the probability of winning amount b, and $1-p-q$ is the probability of winning amount c. The expected value of the lottery is $\mathbf{E}[l_2] = pa + qb + (1-p-q)c$.

A lottery with n payoffs is given in pane (c) of figure 2.1. The prizes are now $a_1, \ldots, a_n$ with probabilities $p_1, \ldots, p_n$, respectively. The expected value of the lottery is now $\mathbf{E}[l_3] = p_1a_1 + p_2a_2 + \ldots + p_na_n$.

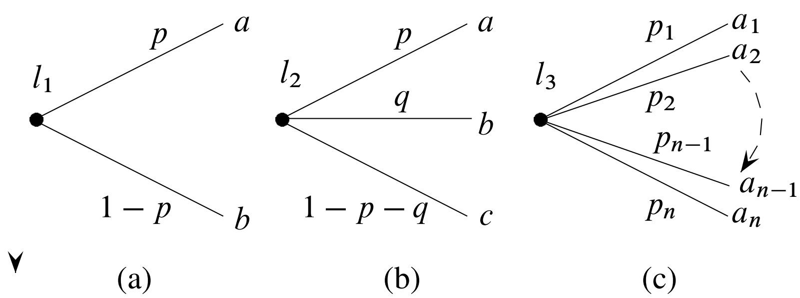

Figure 2.1. Lotteries with two, three, and n potential outcomes

In this section we generalize the previous argument, developing a set of behavioral properties that yield both a utility function over outcomes and a probability distribution over states of nature, such that the expected utility principle, defined in theorem 2.3 holds. Von Neumann and Morgenstern (1944), Friedman and Savage (1948), Savage (1954), and Anscombe and Aumann (1963) showed that the expected utility principle can be derived from the assumption that individuals have consistent preferences over an appropriate set of lotteries. We outline here Savage's classic analysis of this problem.

For the rest of this section, we assume $\succeq$ is a preference relation (§2.1). To ensure that the analysis is not trivial, we also assume that $x \succeq y$ is false for at least some $x, y \in X$. Savage's accomplishment was to show that if the individual has a preference relation over *lotteries* that has some plausible properties, then not only can the individual's preferences be represented by a utility function, but we can infer the probabilities the individual implicitly places on various events, and the expected utility principle (theorem 2.3) holds for these probabilities. These probabilities are called the individuals *subjective prior*.

To see this, let Ω be a finite set of *states of nature*. We call $A \subseteq \Omega$ *events*. Also, let $\mathcal{L}$ be a set of "lotteries," where a *lottery* is a function $\pi : \Omega \to X$ that associates with each state of nature $\omega \in \Omega$ a payoff $\pi(\omega) \in X$. Note that this concept of a lottery does not include a probability distribution over the states of nature. Rather, the Savage axioms allow us to associate a subjective prior over each state of nature ω, expressing the decision maker's personal assessment of the probability that ω will occur. We suppose that the individual chooses among lotteries without knowing the state of nature, after which "Nature" choses the state $\omega \in \Omega$ that obtains, so that if the individual chose lottery $\pi \in \mathcal{L}$, his payoff is $\pi(\omega)$.

Now suppose the individual has a preference relation $\succ$ over $\mathcal{L}$ (we use the same symbol $\succ$ for preferences over both outcomes and lotteries). We seek a set of plausible properties of $\succ$ over lotteries that together allow us to deduce (a) a utility function $u : X \to \mathbf{R}$ corresponding to the preference relation $\succ$ over outcomes in X; (b) a probability distribution $p : \Omega \to \mathbf{R}$ such that the expected utility principle holds with respect to the preference relation $\succ$ over lotteries and the utility function $u(\cdot)$; that is, if we define

$$\mathbf{E}_\pi[u;p] = \sum_{\omega \in \Omega} p(\omega)u(\pi(\omega)), \tag{2.8}$$

then for any $\pi, \rho \in \mathcal{L}$,

$$\pi \succ \rho \Longleftrightarrow \mathbf{E}_\pi[u;p] > \mathbf{E}_\rho[u;p].$$

Our first condition is that $\pi \succ \rho$ depends only on states of nature where π and ρ have different outcomes. We state this more formally as

A1. For any $\pi, \rho, \pi', \rho' \in \mathcal{L}$, let $A = \{\omega \in \Omega | \pi(\omega) \neq \rho(\omega)\}$. Suppose we also have $A = \{\omega \in \Omega | \pi'(\omega) \neq \rho'(\omega)\}$. Suppose also that $\pi(\omega) = \pi'(\omega)$ and $\rho(\omega) = \rho'(\omega)$ for $\omega \in A$. Then $\pi \succ \rho \Leftrightarrow \pi' \succ \rho'$.

This axiom allows us to define a *conditional preference* $\pi \succ_A \rho$, where $A \subseteq \Omega$, which we interpret as "π is strictly preferred to ρ, conditional on event A," as follows. We say $\pi \succ_A \rho$ if, for some $\pi', \rho' \in \mathcal{L}$, $\pi(\omega) = \pi'(\omega)$ and $\rho(\omega) = \rho'(\omega)$ for $\omega \in A$, $\pi'(\omega) = \rho'(\omega)$ for $\omega \notin A$, and $\pi' \succ \rho'$. Because of A1, this is well defined (that is, $\pi \succ_A \rho$ does not depend on the particular $\pi', \rho' \in \mathcal{L}$). This allows us to define $\succeq_A$ and $\sim_A$ in a similar manner. We then define an event $A \subseteq \Omega$ to be *null* if $\pi \sim_A \rho$ for all $\pi, \rho \in \mathcal{L}$.

Our second condition is then the following, where we write $\pi = x|A$ to mean $\pi(\omega) = x$ for all $\omega \in A$ (that is, $\pi = x|A$ means π is a lottery that pays x when A occurs).

A2. If $A \subseteq \Omega$ is not null, then for all $x, y \in X$, $\pi = x|A \succ_A \pi = y|A \Leftrightarrow x \succ y$.

This axiom says that a natural relationship between outcomes and lotteries holds: if π pays x given event A and ρ pays y given event A, and if $x \succ y$, then $\pi \succ_A \rho$, and conversely.

Our third condition asserts that the probability that a state of nature occurs is independent from the outcome one receives when the state occurs. The difficulty in stating this axiom is that the individual cannot choose probabilities, but only lotteries. But, if the individual prefers x to y, and if $A, B \subseteq \Omega$ are events, then the individual treats A as "more probable" than B if and only if a lottery that pays x when A occurs and y when A does not occur will be preferred to a lottery that pays x when B occurs and y when B does not. However, this must be true for any $x, y \in X$ such that $x \succ y$, or the individual's notion of probability is incoherent (that is, it depends on what particular payoffs we are talking about. For instance, some people engage in "wishful thinking," where if the prize associated with an event increases, the individual thinks it is more likely to occur). More formally, we have the following, where we write $\pi = x, y|A$ to mean "$\pi(\omega) = x$ for $\omega \in A$ and $\pi(\omega) = y$ for $\omega \notin A$."

A3. Suppose $x \succ y$, $x' \succ y'$, $\pi, \rho, \pi', \rho' \in \mathcal{L}$, and $A, B \subseteq \Omega$. Suppose that $\pi = x, y|A$, $\rho = x', y'|A$, $\pi' = x, y|B$, $\rho' = x', y'|B$. Then $\pi \succ \pi' \Leftrightarrow \rho \succ \rho'$.

The fourth condition is a weak version of *first-order stochastic dominance*, which says that if one lottery has a higher payoff than another for any event, then the first is preferred to the second.

A4. For any event A, if $x \succ \rho(\omega)$ for all $\omega \in A$, then $\pi = x|A \succ_A \rho$. Also, for any event A, if $\rho(\omega) \succ x$ for all $\omega \in A$, then $\rho \succ_A \pi = x|A$.

In other words, if for any event A, $\pi = x$ on A pays more than the best ρ can pay on A, the $\pi \succ_A \rho$, and conversely.

Finally, we need a technical property to show that a preference relation can be represented by a utility function. It says that for any $\pi, \rho \in \mathcal{L}$, and any $x \in X$, we can *partition* Ω into a number of disjoint subsets $A_1, \ldots A_n$ such that $\cup_i A_i = \Omega$, and for each A_i, if we change π so that its payoff is x on A_i, then π is still preferred to ρ. Similarly, for each A_i, if we change ρ so that its payoff is x on A_i, then π is still preferred to ρ. This means that no payoff is "supergood," so that no matter how unlikely an event A is, a lottery with that payoff when A occurs is always preferred to a lottery with a different payoff when A occurs. Similarly, no payoff can be "superbad." The condition is formally as follows:

A5. For all $\pi, \pi', \rho, \rho' \in \mathcal{L}$ with $\pi \succ \rho$, and for all $x \in X$, there are disjoint subsets $A_1, \ldots, A_n$ of Ω such that $\cup_i A_i = \Omega$ and for any A_i (a) if $\pi'(\omega) = x$ for $\omega \in A_i$ and $\pi'(\omega) = \pi(\omega)$ for $\omega \notin A_i$, then $\pi' \succ \rho$, and (b) if $\rho'(\omega) = x$ for $\omega \in A_i$ and $\rho'(\omega) = \rho(\omega)$ for $\omega \notin A_i$, then $\pi \succ \rho'$.

We then have Savage's theorem.

THEOREM 2.3 *Suppose* **A1–A5** *hold. Then there is a probability function p on Ω and a utility function $u : X \to \mathbf{R}$ such that for any $\pi, \rho \in \mathcal{L}$, $\pi \succ \rho$ if and only if* $\mathbf{E}_\pi[u; p] > \mathbf{E}_\rho[u; p]$.

The proof of this theorem is somewhat tedious (it is sketched in Kreps 1988).

We call the probability p the individual's *Bayesian prior*, or *subjective prior*, and say that A1–A5 imply *Bayesian rationality*, because they together imply Bayesian probability updating.

2.4 Risk and the Shape of the Utility Function

If $\succeq$ is defined over X, we can say nothing about the *shape* of a utility function $u(\cdot)$ representing $\succeq$, because by theorem 2.2, any increasing function of $u(\cdot)$ also represents $\succeq$. However, if $\succeq$ is represented by a utility function $u(x)$ satisfying the expected utility principle, then $u(\cdot)$ is determined up to an arbitrary constant and unit of measure.[4]

THEOREM 2.4 *Suppose the utility function $u(\cdot)$ represents the preference relation $\succeq$ and satisfies the expected utility principle. If $v(\cdot)$ is another utility function representing $\succeq$, then there are constants $a, b \in \mathbf{R}$ with $a > 0$ such that $v(x) = au(x) + b$ for all $x \in X$.*

[4]Because of this theorem, the difference between two utilities means nothing. We thus say utilities over outcomes are *ordinal*, meaning we can say that one bundle is preferred to another, but we cannot say by how much. By contrast, the next theorem shows that utilities over lotteries are *cardinal*, in the sense that, up to an arbitrary constant and an arbitrary positive choice of units, utility is numerically uniquely defined.

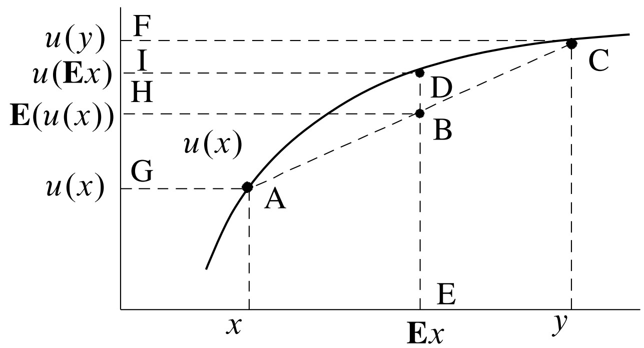

Figure 2.2. A concave utility function

see Mas-Colell, Whinston, and Green (1995):173 prove this theorem.

If $X = \mathbf{R}$, so the payoffs can be considered to be "money," and utility satisfies the expected utility principle, what shape do such utility functions have? It would be nice if they were linear in money, in which case expected utility and expected value would be the same thing (why?). But generally utility will be *strictly concave*, as illustrated in figure 2.2. We say a function $u\colon X \to \mathbf{R}$ is strictly concave, if for any $x, y \in X$, and any $p \in (0, 1)$, we have $pu(x) + (1 - p)u(y) < u(px + (1 - p)y)$. We say $u(x)$ is *weakly concave*, or simply *concave* if, for any $x, y \in X$, $pu(x) + (1 - p)u(y) \leq u(px + (1 - p)y)$.

If we define the lottery π as paying x with probability p and y with probability $1 - p$, then the condition for strict concavity says that *the expected utility of the lottery is less than the utility of the expected value of the lottery*, as depicted in figure 2.2. To see this, note that the expected value of the lottery is $E = px + (1 - p)y$, which divides the line segment between x and y into two segments, the segment xE having length $(px + (1 - p)y) - x = (1 - p)(y - x)$, and the segment Ey having length $y - (px + (1 - p)y) = p(y - x)$. Thus, E divides $[x, y]$ into two segments whose lengths have ratio $(1 - p)/p$. From elementary geometry, it follows that B divides segment $[A, C]$ into two segments whose lengths have the same ratio. By the same reasoning, point H divides segments $[F, G]$ into segments with the same ratio of lengths. This means the point H has the coordinate value $pu(x) + (1 - p)u(y)$, which is the expected utility of the lottery. But by definition, the utility of the expected value of the lottery is at D, which lies above H. This proves that the utility of the expected value is greater than the expected value of the lottery for a strictly concave utility function. This is know as *Jensen's inequality*.