2 Expected utility with known probabilities – "risk" – and unknown utilities

The important financial decisions in our life concern large stakes, and then the maximization of expected value may not be reasonable. Most of our decisions concern nonquantitative outcomes such as health states. Then expected value cannot be used because it cannot even be defined. For these reasons, a more general theory is warranted. We now turn to such a theory – expected utility. For simplicity, we consider only the case where probabilities are known in this chapter. This case is called decision under risk. The general case of both unknown probabilities and unknown utility is more complex, and will be dealt with in later chapters. Whereas Chapter 1 showed how to read the minds (beliefs, i.e. subjective probabilities) of people, this chapter will show how to read their hearts (happiness, i.e. utility).

2.1 Decision under risk as a special case of decision under uncertainty

Probabilities can be (un)known to many degrees, all covered by the general term uncertainty. Decision under risk is the special, limiting case where probabilities are objectively given, known, and commonly agreed upon. Risk is often treated separately from uncertainty in the literature. It is more efficient, and conceptually more appropriate, to treat risk as a special case of uncertainty. I will discuss this point in some detail in this and the following sections. This point will be especially important in the study of ambiguity in Part III. Machina (2004) provided a formal model supporting this point.

Throughout this chapter and in fact throughout this whole book, Structural Assumption 1.2.1 (decision under uncertainty) is maintained. Assumptions will be added. All results from uncertainty can immediately be applied to risk, and results from risk can often be extended to uncertainty. I commonly advise students to work on uncertainty rather than on risk as much as possible, because the numerical probability scale often confounds rather than simplifies what is essential. This point also explains the zigzag structure of Part I of this book. The first big chunk upcoming, Chapters 2 and 3, concerns decision under risk, and the second big chunk, Chapter 4, concerns decision under uncertainty. We nevertheless started in Chapter 1 with decision under uncertainty and not with decision under risk. I did so to encourage the readers, from the beginning, to think of decision under risk as being embedded in decision under uncertainty, and never to have thought of decision under risk in isolation.

We begin with an introductory example that can be skipped by theoretically oriented ("c") readers.

Example 2.1.1 [Vendor with probabilities given]. Assume Example 1.1.1. Assume that, in addition, there is extensive statistical data available about the weather. One out of four times there will be no rain, one out of four times it will be all rain, and half the times there is some rain. This data is objective and agreed upon by everyone. Writing P for objective probability, we have $P(s_1) = P(s_3) = ¼$ and $P(s_2) = ½$.

Prospect x (taking ice cream) generates (¼: 400, ½: 100, ¼: −400), meaning that it yields 400 with probability ¼, 100 with probability ½, and −400 with probability ¼. Similarly, y generates (¼:−400, ½: 100, ¼: 400), 0 generates (1:0) (also denoted as just 0), and x + y generates (½:200, ½:0). Note that y generates the same probability distribution over outcomes as x does, assigning the same probability to each outcome. In decision under risk we assume that the *preference value* of a prospect (i.e., its indifference class) is determined entirely by the probability distribution it generates over the outcomes. Then x and y have the same preference value and are equivalent. For determining preference it then suffices to know the probability distribution generated over the outcomes by a prospect. Hence we often give only that probability distribution. Then x and y are both described by (¼: 400, ½: 100, ¼: −400) and they are equated.

The assumption $x \sim y$ can be violated if utility (defined later) is state dependent. Then the goodness or badness of gaining or losing 400 depends on the weather conditions. We will assume that there is no such dependence. □

For decision under risk, we assume that an objective probability P is given on S, assigning to each event E its probability P(E). Then, with p_j denoting $P(E_j)$, each prospect $(E_1{:}x_1,\ldots,E_n{:}x_n)$ generates a probability distribution $(p_1{:}x_1,\ldots,p_n{:}x_n)$ over the outcomes, assigning probability p_j to each outcome x_j. Probability distributions over outcomes taking only finitely many values are called *probability-contingent prospects*. We now define decision under risk.

Assumption 2.1.2 [*Decision under risk*]. Structural Assumption 1.2.1 (decision under uncertainty) holds. In addition, an objective probability measure P is given on the state space S, assigning to each event E its probability P(E). Different event-contingent prospects that generate the same probability-contingent prospect are preferentially equivalent. □

Because of Assumption 2.1.2, it suffices to describe only the generated probability-contingent prospect for determining the preference value of a prospect, without specifying the underlying event-contingent prospect. This will be the approach of the next section.

Example 2.1.3 [Decision under risk violated]. Assume a small urn with 20 balls numbered 1 to 20, and a large urn with 200 balls numbered 1 to 200, where from each a random ball will be drawn and its number inspected. Many subjects prefer gambling on numbers 1–10 from the large urn – for, let us say, \$100 – to gambling on number 1 from the small urn ("I have more chances"), even though both event-contingent prospects generate the same probability-contingent prospect (1/20:100, 19/20:0) (Kirkpatrick & Epstein 1992). This implies an empirical violation of Assumption 2.1.2. □

Exercise 2.1.1.[a] Reconsider Exercise 1.6.7, with Paul and everyone agreeing that the $P(E_j)$'s = ⅙ are objective probabilities. Does Assumption 2.1.2 hold for Paul? □

The most commonly used measure on the real axis assigns to each interval $[\alpha,\beta]$ its length $\beta - \alpha$, and to disjoint unions of intervals the sum of their separate lengths. It is called the *Lebesgue measure*, often denoted λ, and we will use it sometimes. On the unit interval [0,1) it is a probability measure.

Exercise 2.1.2.[a] Assume weak ordering, and Assumption 2.1.2.

(a) Assume that S = [0,1) and that the probability measure P on S is the Lebesgue measure. Let x be the event-contingent prospect ([0,⅓):3, [⅓,⅔):8, [⅔,1):2). What probability-contingent prospect is generated by x? Let y be the event-contingent prospect ([0,⅓):3, [⅓,⅔):4, [⅔,1):5). What probability-contingent prospect is generated by y? Assume that y $\succcurlyeq$ x. Let f be the event-contingent prospect ([0,⅓):8, [⅓,⅔):2, [⅔,1):3). Let g be the event-contingent prospect ([0,⅓):5, [⅓,⅔):4, [⅔,1):3). What is the preference between f and g?

(b) Consider the probability-contingent prospect (¾:8,¼:2). Define two different event-contingent prospects that generate it. □

Example 2.1.4 [Regret (Bardsley *et al.* 2010 Box 3.5)]. Violations of Assumption 2.1.2 may arise because of regret, where you may care about correlations between outcomes (Bell 1982). Assume that S = [0,1) is endowed with the Lebesgue measure as objective probability P, reflecting a random choice of a number from [0,1). Consider a preference

$$([0,\tfrac{1}{6}):10,[\tfrac{1}{6},\tfrac{2}{6}):20,\ [\tfrac{2}{6},\tfrac{3}{6}):30,[\tfrac{3}{6},\tfrac{4}{6}):40,[\tfrac{4}{6},\tfrac{5}{6}):50,[\tfrac{5}{6},1):60) \prec$$

$$([0,\tfrac{1}{6}):60,[\tfrac{1}{6},\tfrac{2}{6}):10,\ [\tfrac{2}{6},\tfrac{3}{6}):20,[\tfrac{3}{6},\tfrac{4}{6}):30,[\tfrac{4}{6},\tfrac{5}{6}):40,[\tfrac{5}{6},1):50).$$

Fishburn (1988 p. 273) and Loomes & Sugden (1982 Table 6) argued that this preference can be rational because the regret for an upper choice under the event of $s<\tfrac{1}{6}$ will be very big, whereas there is only low regret for the lower prospect under all other events. However, the two event-contingent prospects generate the same probability-contingent prospect, assigning probability ⅙ to each outcome. Hence they must be equivalent under the assumption of decision under risk, and the preference violates decision under risk. As the aforementioned references show, regret leads to violations of transitivity. We will throughout assume Assumption 2.1.2 of decision under risk, and will not consider regret or violations of transitivity. □

2.2 Decision under risk: basic concepts

In decision under risk, as in decision under uncertainty, the *outcome set* is $\mathbb{R}$, with real numbers designating money. As defined in the preceding section, *probability-contingent prospects*, or, briefly, *prospects*, are probability distributions over $\mathbb{R}$ that take only finitely many values.[1] The generic notation for a prospect is $(p_1{:}x_1, \ldots, p_n{:} x_n)$, yielding outcome x_j with probability p_j for each j. Here n is a natural number that can be different for different prospects. Again, we often drop brackets, commas, and colons, writing $p_1x_1\ldots p_nx_n$ if no confusion with multiplication is likely to arise; another notation is p_1x_1,p_2x_2,p_3x_3. If there are only two outcomes and probabilities, then we often suppress the second probability and write $x_{1p_1}x_2$ instead of $(p_1{:}x_1, p_2{:} x_2)$, and $\alpha_p\beta$ instead of $(p{:}\alpha, 1-p{:}\beta)$. For a given prospect, the probabilities p_j are the *outcome probabilities*. It is implicitly understood that outcome probabilities are nonnegative and sum to 1. A prospect x, being a probability distribution over the reals, assigns to each interval in the reals the probability that the outcome of the prospect will fall into that interval. The *preference relation* over prospects is again denoted by $\succcurlyeq$.

The figure illustrates the prospect $(\frac{3}{6}{:}20, \frac{1}{6}{:}14, \frac{2}{6}{:}12)$. The circle at the beginning is a *chance node*. *Branches* lead to the outcomes. The corresponding probabilities have been indicated.

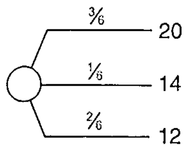

The next figure illustrates the riskless prospect (1:10). In general, a *riskless*, or *degenerate*, prospect yields one fixed outcome with probability 1. It is identified with the outcome. For instance, the outcome 10 is identified with the depicted riskless prospect.

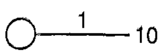

The following assumption is commonly made for decision under risk, and we make it too. It requires that the underlying state space that generates the randomness is sufficiently rich to generate all probabilities. In particular, S is infinite.

Assumption 2.2.1 [*Richness for decision under risk*]. Every probability distribution over the outcomes that takes only finitely many values is available in the preference domain. □

[1] Formally: There is a finite set of outcomes that has probability 1.

Although probabilities are unknown in most practical decision situations, cases of known probabilities are still important. In such cases decision theory can be applied especially fruitfully, in particular if several persons are involved so that agreement and clear communication about probabilities are desirable. Decisions under risk can also serve as a useful benchmark for general decisions under uncertainty, for example when defining attitudes towards ambiguity, as will be explained later.

2.3 Decision under risk as a special case of decision under uncertainty; continued

This is one of several sections in this book that can be skipped without loss of continuity. Such possibilities can be inferred from Appendix K. Hence, in similar cases in future sections this point will not be mentioned again.

It is useful to keep in mind that decision under risk is a special case of decision under uncertainty with the underlying state space suppressed for the sake of convenience. The state space describes the physical process that generates the randomness underlying the probabilities. The following example shows that for decision under risk without a state space specified, a state space (we will construct $S = [0,1)$) can always be defined to serve as underlying state space after all, with respect to which Structural Assumption 1.2.1 (decision under uncertainty) is satisfied.

Example 2.3.1. We assume decision under uncertainty with the unit interval [0,1) as (infinite) state space S, as in Example 2.1.4. In what follows, it will be notationally convenient to have S and other events left-closed and right-open (so that disjoint unions generate same kinds of events). We imagine that an arbitrary number will be drawn from S, and the true state s is the number that will be drawn. Event $[\frac{3}{6}, \frac{4}{6})$, for instance, refers to the event that the number drawn will weakly exceed $\frac{3}{6}$ but will be strictly below $\frac{4}{6}$. An example of an event-contingent prospect is $([0, \frac{3}{6}) : 20, [\frac{3}{6}, \frac{4}{6}) : 14, [\frac{4}{6}, 1) : 12)$, yielding outcome 20 if the number s drawn is less than $\frac{3}{6}$, yielding outcome 14 if $\frac{3}{6} \le s < \frac{4}{6}$, and yielding outcome 12 if $\frac{4}{6} \le s < 1$. We will henceforth use event-contingent prospects – random variables – defined on intervals, such as $x = ([0,q_1):x_1, [q_1,q_2):x_2, \ldots, [q_{n-1},1):x_n)$, yielding x_j if $q_{j-1} \le s < q_j$, for $0 = q_0 < q_1 < \cdots < q_n = 1$.

We use the Lebesgue measure λ defined in §2.1 as probability measure on $S = [0,1)$. Writing $p_j = q_j - q_{j-1}$, the prospect x yields outcome x_1 with probability $p_1, \ldots,$ and outcome x_n with probability p_n. It constitutes a probability distribution $p_1x_1 \ldots p_nx_n$ over the reals, with the p_j's the outcome probabilities. In other words, it generates a probability-contingent prospect. The event-contingent prospect $([0, \frac{3}{6}) : 20, [\frac{3}{6}, \frac{4}{6}) : 14, [\frac{4}{6}, 1) : 12)$, for example, generates the depicted probability-contingent prospect $(\frac{3}{6} : 20, \frac{1}{6} : 14, \frac{2}{6} : 12)$. The event-contingent prospects $([0, \frac{1}{3}) : 30, [\frac{1}{3}, \frac{2}{3}) : 20, [\frac{2}{3}, 1) : 10)$ and $([0, \frac{1}{3}) : 20, [\frac{1}{3}, \frac{2}{3}) : 10, [\frac{2}{3}, 1) : 30)$ both yield the outcomes 10, 20, or 30, each with probability $\frac{1}{3}$, so that both generate the probability-contingent prospect ($\frac{1}{3}: 10, \frac{1}{3}: 20, \frac{1}{3}: 30$). We assume Assumption 2.1.2. Assumptions 1.2.1 and 2.2.1 are also satisfied. □

The following example illustrates a case where both subjective and objective probabilities are given, and where they must agree. For simplicity, we assume expected value in the example. That the result holds in almost complete generality, also if expected value is not assumed, is shown after.

Example 2.3.2 [Objective and subjective probabilities agree under expected value maximization]. Assume Example 2.1.1. Further assume that decision under risk holds, and that the conditions of de Finetti's Theorem 1.6.1 also hold. Then preferences maximize expected value with respect to a subjective probability measure Q that in principle might be different from P. We show

$$Q = P. \tag{2.3.1}$$

Explanation. The prospects (s_1:0, s_2:100, s_3:0) and (s_1:100, s_2:0, s_3:100) generate the same objective probability distribution over outcomes, namely $100_{\frac{1}{2}}0$. Hence, they are equivalent by the assumption of decision under risk. This implies that their subjective expected values, $Q(s_2)100$ and $(Q(s_1) + Q(s_3))100$, must be the same. Hence, $Q(s_2) = (Q(s_1) + Q(s_3)) = \frac{1}{2}$. The prospects ($s_1$:100, s_2:0, s_3:0) and (s_1:0, s_2:0, s_3:100) also generate the same objective probability distribution over outcomes, namely ($\frac{1}{4}$:100, $\frac{3}{4}$:0), and are therefore also equivalent by the assumption of decision under risk. Then $Q(s_1)100 = Q(s_3)100$, and $Q(s_1) = Q(s_3)$ follows. Because these subjective probabilities sum to $\frac{1}{2}$, each must be $\frac{1}{4}$. We conclude that Eq. (2.3.1) holds. Hence, preferences maximize objective expected value in this example. □

The following exercise shows that, under mild assumptions, subjective probabilities must agree with objective probabilities, also if expected value maximization does not hold. This result will hold for all models considered later in this book. The result implies that the (subjective) probability models developed by de Finetti (1931a), Savage (1954), and others do not *deviate* from objective probability models, but *generalize* them by also incorporating situations in which no objective probabilities are given. There have been some misunderstandings about this issue, primarily in the psychological literature. It was sometimes believed, erroneously, that the ("subjective") expected utility models of Savage (1954) and others were developed to allow subjective probabilities to deviate from objective probabilities (Edwards 1954 pp. 396–397, corrected in Edwards 1962 p. 115; Lopes 1987 p. 258). In modern papers, the term subjective probability is still sometimes used in this unfortunate sense. This misunderstanding may have contributed to the unfortunate separation between risk and uncertainty in some parts of the literature. Mosteller & Nogee (1951 p. 398 and footnote 16) presented a clear and correct discussion.

Exercise 2.3.1^b [Subjective probabilities must agree with objective probabilities under stochastic dominance]. Assume that $S = [0,1)$ with P the Lebesgue measure ($P([a,b)) = b - a$), and assume weak ordering. Define $E_j = [\frac{j-1}{6}, \frac{j}{6})$. Consider event-contingent prospects of the form $E_1x_1...E_6x_6$, assigning x_j to all states s in the interval $[\frac{j-1}{6}, \frac{j}{6})$. Assume (deviating from Exercise 1.6.7) that decision under risk holds in two ways at the same time for the preference relation $\succcurlyeq$: it holds with respect to the objective probability measure P, but also with respect to a, possibly different, subjective probability measure Q. Note that this concerns one and the same preference relation, for which both of these assumptions hold at the same time.

Regarding P, the value of a prospect depends only on the probability distribution generated over outcomes through P. For example, the event-contingent prospect $100_{E1}0$, yielding 100 if the true state is below $\frac{1}{6}$ and yielding zero otherwise, is equivalent to the event-contingent prospect $100_{E6}0$, yielding 100 if the true state is at least $\frac{5}{6}$ and yielding 0 otherwise. We further assume a strict stochastic dominance condition with respect to Q: $100_A0 \succ 100_B0$ whenever $Q(A) > Q(B)$. This condition holds, for instance, under EV maximization. Write $q_j = Q(E_j)$. Show that $q_1 = \cdots = q_6 = \frac{1}{6}$.

Comment.' For any event E with objective probability $P(E) = \frac{1}{6}$ we can construct a 6-fold partition of S with E one of the elements and all elements of the partition having P-value $\frac{1}{6}$. It follows in the same way as in Exercise 2.3.1 that these events have Q-value $\frac{1}{6}$, too. We can in the same way show that all events with objective probability 1/m must also have subjective probability 1/m. By taking unions we see that all events with objective rational probabilities j/m have the same subjective probabilities j/m. It then readily follows (bounding between rational-probability sub- and supersets) that irrational (in the mathematical sense of being no fraction of integers) objective probabilities must also be identical to subjective probabilities. Hence, for rich event spaces, subjective probabilities must indeed agree with objective probabilities.

Further Comment. I briefly discuss situations in which the result of this exercise can be violated. One assumption underlying our analysis is typical of decision theory and is made throughout this book: All prospects are conceivable and the preference relation is complete over all these conceivable prospects. De Finetti, more generally, proved his bookmaking theorem (Theorem 1.6.1) on any linear subspace of prospects. Such cases are important in finance. Financial markets are often incomplete, i.e. it cannot be assumed that all conceivable prospects are available. For example, a prospect 100_E0 with E the event that the economy goes down dramatically and every one is losing is not plausible. It may then happen that the result of this exercise does not hold, and risk neutral market ("subjective") probabilities can deviate from objective statistical probabilities. The deviation may be generated by risk aversion and state dependence of utility (see Example 2.1.1; it will be defined formally later) for instance, with different marginal utilities for different events. This point is central to many works (Kadane & Winkler 1988; Karni 1996; Nau 2003 opening paragraph). The discrepancy between market probabilities and objective probabilities can be used to estimate risk aversion (Bliss & Panigirtzoglou 2004).[2] □

[2] Pfanzagl (1968 Theorem 12.6.17) presented further results on the equality of objective and subjective probabilities.

Assignment 2.3.2.[b] Assume weak ordering, monotonicity, and Assumption 2.1.2 (decision under risk). Assume $S = [0,1)$, with the probability measure P on S the Lebesgue measure. Let x be the event-contingent prospect $([0,\frac{1}{3}){:}1, [\frac{1}{3},\frac{2}{3}){:}2, [\frac{2}{3},1){:}3)$, and let y be the event-contingent prospect $([0,\frac{1}{3}){:}3, [\frac{1}{3},\frac{2}{3}){:}4, [\frac{2}{3},1){:}2)$. Show that $y \succ x$. □

2.4 Choices under risk and decision trees

Fig. 2.4.1a shows two prospects, the upper one yielding \$80 with probability 0.5 and \$60 with probability 0.5, and the lower one yielding \$95 with probability 0.5 and \$70 with probability 0.5. The figure displays a choice situation in which you have to choose between the upper and the lower prospect. The square is a *decision node*; that is, it indicates a situation in which you must choose. Such figures are read from left to right. For instance, Fig. a illustrates a situation where first you arrive at the square node and have to choose if you prefer up or down. If you choose up, you end up in the circle designating the prospect $80_{\frac{1}{2}}60$. If you choose down, you receive $95_{\frac{1}{2}}70$. Fig. b again starts with a decision node where you are to decide which way to go, up or down. Then chance nodes result where chance decides where to go. In Fig. b you have to choose between two prospects each yielding \$40 or \$20, yielding the higher outcome with probability 0.7 (after up in the decision node) or with probability 0.6 (after down in the decision node).

Exercise 2.4.1.[a] Determine in each of the eight situations depicted in Figure 2.4.1 whether you prefer the upper or the lower prospect. Keep your answers for later use. □

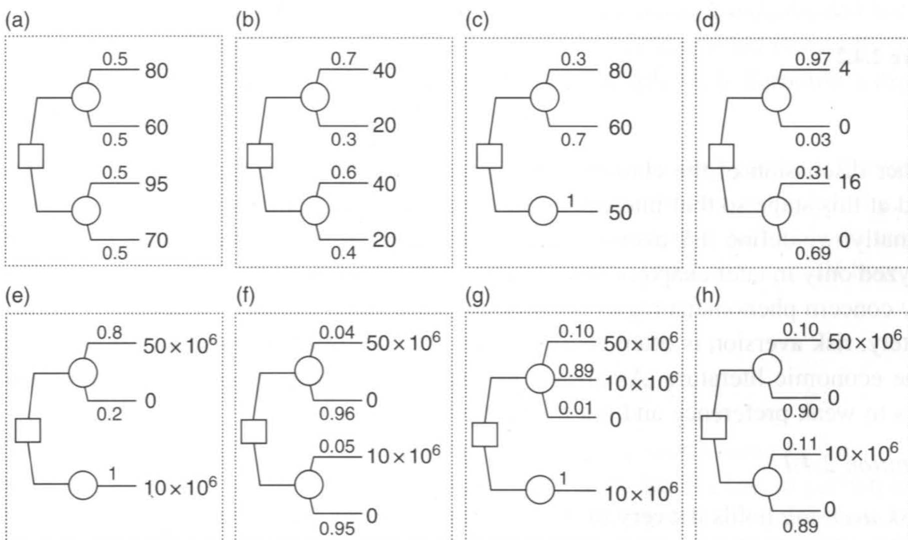

Figure 2.4.1

The choice situations in (a), (b), and (c) were trivial. The lower prospect in (a) results from the upper one by replacing the outcome 80 by the outcome 95 and the outcome 60 by the outcome 70. In other words, the lower prospect simply gives better outcomes. It is therefore obvious that the lower prospect is preferable. In (b), the upper prospect results from the lower one by shifting 0.1 probability mass from the bad outcome 20 to the good outcome 40, which again is a definite improvement. In (c), the lower prospect results from the upper one by replacing both outcomes by the worse outcome 50 (or, by shifting all probability mass from the outcomes 80 and 60 to the worse outcome 50) and, hence, it is obvious that the upper prospect is preferable.

If you are indifferent, in Fig. 2.4.1e, between the upper and the lower prospect, then your certainty equivalent for the upper prospect is $\$10 \times 10^6$. It implies that you would equally well like to receive $\$10 \times 10^6$ for sure as $\$50 \times 10^6$ with probability 0.8 (and nothing with probability 0.2).

Exercise 2.4.2.[a] Substitute your own certainty equivalent for both prospects depicted in Figure 2.4.2.

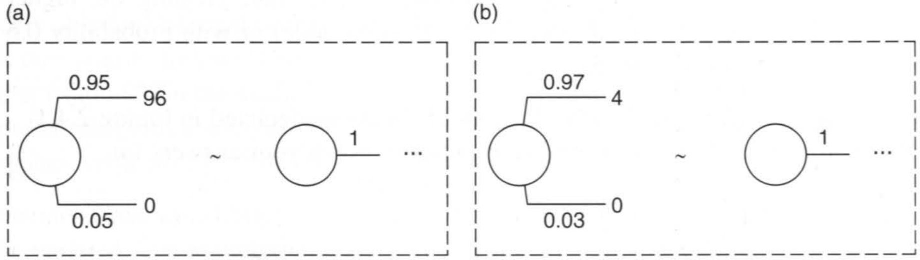

Figure 2.4.2

□

Further discussion of the choices is postponed until §2.8, although it can be understood at this stage so that interested readers can read §2.8 immediately.

Finally, we define risk aversion and risk seeking. Although these concepts will be analyzed only in later chapters, they are defined here to stress their theory-free status. They concern phenomena regarding preferences without any theory assumed. Unfortunately, risk aversion is often confused with concave utility (concepts defined later) in the economic literature. As always in this book, preference without qualification refers to weak preference and not to strict preference.

Definition 2.4.1.

- *Risk aversion* holds if every prospect is less preferred than its expected value.
- *Risk seeking* holds if every prospect is preferred to its expected value.
- *Risk neutrality* holds if every prospect is indifferent to its expected value. □

Risk aversion implies, for example, $50 \succcurlyeq 100_{½}0$, and risk seeking implies $50 \preccurlyeq 100_{½}0$. Risk neutrality means that both risk aversion and risk seeking hold, and it is just another name for expected value maximization. Then $50 \sim 100_{½}0$. It was used in Chapter 1.

2.5 Expected utility and utility measurement

De Finetti's Theorem 1.6.1 gave a behavioral foundation for expected value maximization for decision under uncertainty. The result can be applied to decision under risk as well. Exercise 2.3.1 and its comment showed that objective and subjective probabilities usually agree. Hence, Theorem 1.6.1, when stated for decision under risk, can serve as a behavioral foundation for expected value maximization.

For expected value maximization under risk there are no free subjective parameters in the model. Preferences are uniquely determined and all expected value maximizers behave the same way. With probabilities and outcomes given, we can immediately calculate expected value. The expected value criterion is thus directly observable by itself and does not need a behavioral foundation. It is its own, trivial, behavioral foundation so to speak. In terms of expected utility for uncertainty defined later, expected value for decision under risk concerns the special case where both probabilities and utilities are known, so that nothing about decision attitude remains to be discovered.

Table 1.5.4 of finance with bankruptcy at −70K illustrated a case where outcomes were so large that violations of expected value were plausible. The rest of this chapter examines this case for decision under risk. Henceforth, probabilities are known but now utilities will be unknown. The following illustration of a violation of expected value is not very realistic and has a didactical confound (the complexities of infinity). To my surprise, experience has shown that it still works well didactically because students like it. The example is historically important because it led Bernoulli (1738) to develop expected utility. We will return to the example with Bernoulli's explanation later.

Example 2.5.1 [St. Petersburg paradox]. Consider the following game. A fair coin will be flipped until the first heads shows up. If heads shows up at the kth flip, then you receive $\$2^k$. Thus, immediate heads gives only \$2, and after each tails the amount doubles. After 19 tails you are sure to be a millionaire. Think for yourself how much it would be worth to you to play this game.
The expected value of the game is

$$½ \times 2 + ¼ \times 4 + ⅛ \times 8 + \tfrac{1}{16} \times 16 + \cdots = 1 + 1 + 1 + 1 + \cdots = \infty.$$

Thus if you maximize expected value, then this game is worth more to you than any amount of money. In reality, people pay considerably less to participate in the game, something like \$5 (as Bernoulli wrote based on casual observations), showing that expected value does not hold empirically when large amounts are involved. □

In Chapter 1 we first introduced a number of plausible preference conditions. We then came to the, possibly surprising, conclusion that a specific preference model was implied, namely EV. §1.8 further discussed general features of such behavioral foundations. We could present all behavioral foundations in the same spirit, especially those with a normative status. For brevity's sake, we will not do so. We follow a different system henceforth, as stated in the Introduction. It is shorter and can be applied both normatively and descriptively, and is as follows. First an empirical measurement of the subjective parameters is described that relates them directly to preferences. Next, consistency of measurement gives a behavioral foundation.

We summarize the structural assumption needed for the following analysis. It is implied by Assumptions 1.2.1 (decision under uncertainty), 2.1.2 (decision under risk), and 2.2.1 (richness), as can be seen.

Structural Assumption 2.5.2 [*Decision under risk and richness*]. $\succcurlyeq$ is a preference relation over the set of all (probability-contingent) prospects, i.e., over all finite probability distributions over the outcome set $\mathbb{R}$. □

Definition 2.5.3. Under Structural Assumption 2.5.2, *expected utility* (*EU*) holds if there exists a strictly increasing function U, the *utility* (*function*), from the outcome set to $\mathbb{R}$, such that the evaluation depicted in Figure 2.5.1 represents preferences. □

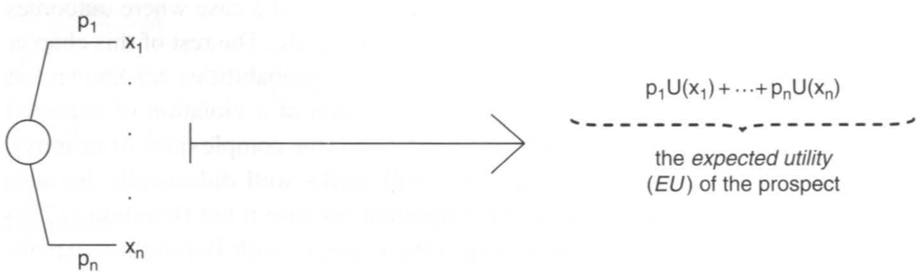

Figure 2.5.1

Under EU, the utility function is the subjective parameter characterizing the decision maker. The following exercise shows how to derive decisions from the EU model. Its elaboration is immediately given in the main text.

Exercise 2.5.1.[a] Assume EU with U(0) = 0, U(100) = 1, U(81) = 0.9. Determine the preference between $100_{0.9}0$ and $100_{0.1}81$. □

In Exercise 2.5.1

$EU(100_{0.9}0) = 0.90 \times U(100) + 0.10 \times U(0) = 0.90$;

$EU(100_{0.1}81) = 0.10 \times U(100) + 0.90 \times U(81) = 0.10 + 0.81 = 0.91$.

Hence $100_{0.1}81$ has the higher EU and is preferred.

Exercise 2.5.2.[a] Assume EU with $U(\alpha) = \sqrt{\alpha}$. Determine the preference between $49_{0.9}16$ and $81_{0.7}16$. □

Example 2.5.4 [St. Petersburg paradox explained using EU]. Consider Example 2.5.1. Bernoulli argued that expected utility with a logarithmic utility function is plausible. Then the expected utility of the game is

$$\sum_{j=1}^{\infty} 2^{-j} ln(2^j) = \sum_{j=1}^{\infty} j\ 2^{-j} ln(2) = 2ln(2) = ln(4).^3$$

Hence the certainty equivalent is \$4, fitting Bernoulli's empirical claims well. □

As in §1.3, we can use the elicitation method, and derive properties of utility from observed choices. Part (a) of the following exercise, and what follows, illustrate this point. The exercise illustrates the applicability of expected utility, with some choices used to assess the model and this assessment subsequently used to predict new choices. Its elaboration is given immediately following the exercise.

Exercise 2.5.3.[a] Assume that you observe some preferences of an agent. You know that the agent satisfies EU, and can set U(0) = 0 and U(100) = 1.

(a) Assume that $50 \sim 100_{0.58}0$ for the agent. What is U(50)?
(b) Under the indifference in part (a), you can predict the agent's preference between (0.40:100, 0.20:50, 0.40:0) and (0.33:100, 0.33:50, 0.34:0). What is it?

Elaboration.

(a) Because of the equivalence, $U(50) = EU(100_{0.58}0) = 0.58 \times 1 + 0.42 \times 0 = 0.58$.
(b) $EU(0.40{:}100, 0.20{:}50, 0.40{:}0) = 0.40 \times 1 + 0.20 \times 0.58 + 0.40 \times 0 = 0.40 + 0.116 + 0 = 0.516$.

$EU(0.33{:}100, 0.33{:}50, 0.34{:}0) = 0.33 \times 1 + 0.33 \times 0.58 + 0.34 \times 0 = 0.33 + 0.191 + 0 = 0.521$.

The latter prospect has the higher EU and is preferred. □

Exercise 2.5.4.[a] Assume EU with U(0) = 0 and U(100) = 1.

(a) Assume that $60 \sim 100_{0.70}0$ for the agent. What is U(60)?
(b) Under the indifference in part (a), what is the preference between $60_{0.70}0$ and $100_{0.49}0$? □

The best way to learn about the general way to measure utility is to invent it by yourself. You can do so in the following exercise. Its elaboration is given next in the main text.

Exercise 2.5.5.[a] Assume EU with U(0) = 0 and U(100) = 1. What kind of preferences should you observe to find out what U(30) and U(70) are? □

[3] Second "=" : $\sum_{j=1}^{\infty} j2^{-j} = \sum_{j=1}^{\infty} 2^{-j} + \sum_{j=1}^{\infty}(j-1)2^{-j} = 1 + \frac{1}{2}\sum_{j=2}^{\infty}(j-1)2^{-(j-1)} = 1 + \frac{1}{2}\sum_{j=1}^{\infty} j2^{-j}$. Then $\sum_{j=1}^{\infty} j2^{-j} = 2$.

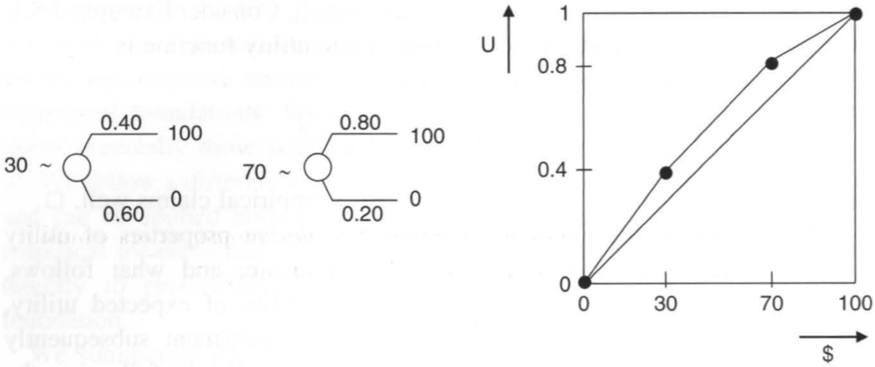

Figure 2.5.2 Two indifferences and the resulting U curve.

Measuring utility. To analyze risky decisions, we want to measure utility U. Figure 2.5.2 shows an example. Assume EU with the normalization $U(0) = 0$ and $U(100) = 1$.

The first indifference implies

$$U(30) = 0.40 \times U(100) + 0.60 \times U(0) = 0.40.$$

The second indifference implies

$$U(70) = 0.80 \times U(100) + 0.20 \times U(0) = 0.80.$$

The graph in Figure 2.5.2 then results from linear interpolation.

In general, assume two fixed outcomes $M > m$, and assume that we have normalized $U(m) = 0$ and $U(M) = 1$. (In Figure 2.5.2, $m = 0$ and $M = 100$.) We will see later (Exercise 2.6.4) that we may always assume such normalized utility values. For each outcome α such that $M \geq \alpha \geq m$, we can elicit the *standard-gamble (SG) probability* with respect to m and M, being the probability p such that the equivalence in Figure 2.5.3 holds. Applying EU gives

$$U(\alpha) = pU(M) + (1 - p)U(m) = p \text{ [\textit{the SG Eq.}]} \tag{2.5.1}$$

In other words, $U(\alpha)$ is simply the SG probability p. Exercises 2.5.3a and 2.5.4a illustrated this procedure.

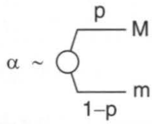

Figure 2.5.3 The SG probability p of α.

The SG method directly relates utility to decisions, in a very simple manner. The method shows how to make exchanges between outcomes and probability. This way outcomes and probabilities become commensurable. Through a SG consistency condition defined later, this commensurability will capture the essence of EU. Surveys of utility measurement under EU are in Farquhar (1984) and Seidl & Traub (1999).

Assignment 2.5.6.[c] This assignment does not concern binary choice, as does most of this book, but it concerns choosing from a continuum of prospects.

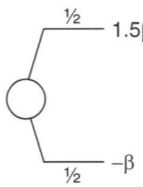

(a) Assume that you can invest an amount β at your discretion in a stock that with ½ probability yields a return of 2.5β and, thus, a profit of 1.5β. There is a probability ½ of ruin, in which case you lose your investment β. You maximize EU, with utility function $U(\alpha) = 1 - e^{-\alpha}$. How much do you invest? What are the resulting EU and CE?
(b) The same question, but now $U(\alpha) = 1 - e^{-\theta\alpha}$ for a general $\theta > 0$. (Hint: substitute $\theta\alpha$ for α in part (a)). □

2.6 Consistency of measurement and a behavioral foundation of expected utility

We assume in this chapter, as we do throughout this book, that the outcome set is $\mathbb{R}$. It can be seen, though, that the behavioral foundation of expected utility presented in Theorem 2.6.3, and the analysis of this whole chapter, hold for general outcome sets, and that we never need the outcome set to be $\mathbb{R}$. The only change required for general outcomes is that inequalities between outcomes such as $M > m$ be replaced by preferences between outcomes such as $M \succ m$.

A necessary condition for EU to hold, and for the measurement of utility to work well, is that a probability p to solve the indifference in Figure 2.5.3 exists for all $M > \alpha > m$. This requirement is called *standard-gamble (SG) solvability*. It is the only continuity-like condition needed in the analysis to follow. It entails that the probability scale is at least as refined as the outcome scale. Solvability conditions are less restrictive and better capture the empirical meaning of a model than continuity (which essentially needs connectedness of the reals in preference foundations; Luce *et al.* 1990 p. 49 l. 10; Wakker 1988).

Another immediate implication of EU, based on utility being strictly increasing, is as follows.

Standard-gamble (SG) dominance. For all outcomes M > m and probabilities p > q:

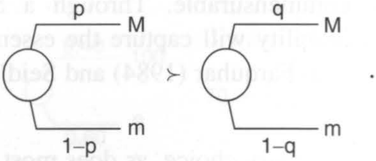

Moving positive probability mass from an outcome to a strictly better outcome entails a strict improvement. The condition is a weak version of stochastic dominance (defined later).

Exercise 2.6.1.[a] Assume EU. Verify, for general U(M) > U(m), that p = (U(α)−U(m))/(U(M)−U(m)) in Figure 2.5.3. □

Exercise 2.6.2.[b] Show that SG dominance holds under EU. In other words, assume a preference relation over prospects that can be represented by EU. Show that the preference relation satisfies SG dominance. □

Exercise 2.6.3.[b] Assume weak ordering, SG dominance, and SG solvability. Show, for M > m, that for the SG probability p in Figure 2.5.3:

(a) p is uniquely determined.
(b) If α ∼ m, then p = 0.
(c) If α ∼ M, then p = 1.
(d) If m < α < M, then 0 < p < 1. □

The following exercise deals with an important property of the utility function in expected utility. If you have no clue how to resolve the exercise, then a comment similar to the one preceding Exercise 1.3.4 applies. You probably have not yet fully internalized the conceptual meaning of decision models. A rethinking of Definition 2.5.3 and Exercise 2.5.3, and a rereading of some of the preceding text, may then be useful. Alchian (1953) gives a didactic and detailed account of the empirical meaning of utility and of the following exercise.

Exercise 2.6.4[!a] [Free level and unit of U]. Show that, if EU holds with a utility function U – i.e., the expectations of U as in Figure 2.5.1 represent preferences – then so does EU with utility function $U^* = \tau + \sigma U$, where $\tau \in \mathbb{R}$ and $\sigma > 0$. □

Exercise 2.6.4 shows that we can normalize the utility scale as we like, and choose where U will be 0 and 1.

Exercise 2.6.5.[c] Assume EU. Show that U is continuous if and only if there exists a certainty equivalent for each prospect. □

The implications of EU just described are relatively uncontroversial. We now turn to the most interesting, and most controversial, implication – the well-known independence condition. We consider here a version called standard gamble consistency, which states the condition in terms of consistency of utility measurement. Most normative discussions of EU concern the independence condition. Empirical tests of EU and its alternatives usually examine how independence is violated, and most nonexpected utility models weaken independence in specific ways. Such alternative models will be presented later.

In a perfect world where EU would hold perfectly well, the utility measurement through the SG probability in Figure 2.5.3 would provide exact measurements of U that would give perfect predictions. Unfortunately, there usually are measurement errors and biases, and models do not hold perfectly well. McCord & de Neufville (1986) criticized the SG method for being particularly prone to empirical biases. The main reason is that this method involves a riskless outcome α, and such outcomes are often evaluated in deviating manners, generating violations of EU. It may, therefore, be safer to avoid their use. McCord & de Neufville, and several other authors (Carthy *et al.* 1999 §2; Davidson, Suppes, & Siegel 1957; Officer & Halter 1968 p. 259), recommended embedding SG indifferences as in Figure 2.5.3 into more complex risky choice prospects. To prepare for this embedding, we discuss the (probabilistic) mixing of prospects such as $x_\lambda y$, clarified in the next example.

Example 2.6.1 [Probabilistic mixing of prospects]. Consider the left-hand two-stage prospect in Figure 2.6.1. In the first chance node there is a ½ probability of going up and a ½ probability of going down. If up, the prospect $200_{1/3}0$ results. These ⅓ and ⅔ probabilities are conditional probabilities, being conditional on the event of going up. The other probabilities are similar. The probability of receiving \$200 is, therefore, $\frac{1}{2} \times \frac{1}{3} = \frac{1}{6}$, the probability of \$100 is also $\frac{1}{6}$, and the probability of 0 is $\frac{1}{2} \times \frac{2}{3} + \frac{1}{2} \times \frac{2}{3} = \frac{2}{3}$. That is, the probability distribution is the same as in the right-hand prospect.

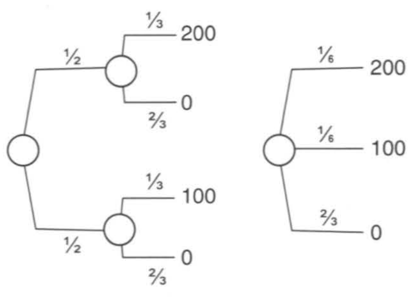

Figure 2.6.1

□

Definition 2.6.2. For general prospects x, y, and $0 \leq \lambda \leq 1$, the mixture depicted in Fig. 2.6.2a, also denoted $x_\lambda y$, is the prospect that assigns, to each outcome α, λ times the probability of α under x plus $1-\lambda$ times the probability of α under y. It is called a (*probabilistic*) *mixture* of x and y. □

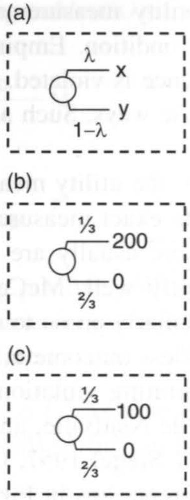

Figure 2.6.2

In Figure 2.6.1, x was the prospect in Fig. 2.6.2b, y was the one in Fig. 2.6.2c, and λ was ½. The mixture $x_\lambda y$ can be thought of as resulting from a two-stage procedure, where in the first stage prospect x results with probability λ and prospect y with probability $1-\lambda$, and in the second stage the prospect resulting from the first stage (x or y, as the case may be) is played out.[4] Figure 2.6.3 shows another example.

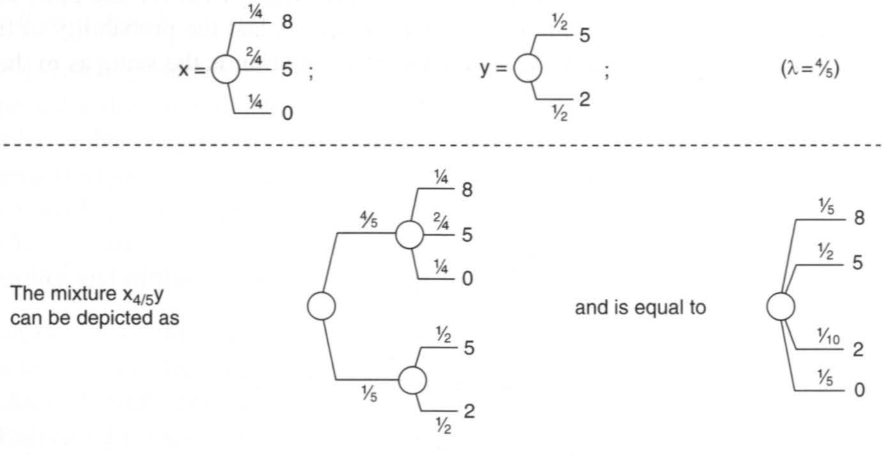

Figure 2.6.3

[4] This footnote is meant to be read only by specialists in decision theory. It concerns the reduction of compound lotteries (= prospects) assumption. If mixtures of prospects are interpreted as physically resulting from a two-stage procedure as just described, with physical time proceeding, then the *reduction of compound prospects assumption* is used, which entails that probabilities of subsequent chance nodes are multiplied. In our analysis, the only objects considered are one-stage prospects. Two-stage illustrations are only a didactic device for depicting one-stage prospects.

I next discuss linearity of EU in probabilities. Because this property is important for understanding EU, I shall explain it in some detail. We first consider this property for general integrals. If with probability ½ you receive a prospect with expected return \$100, and with probability ½ you receive a prospect with expected return \$200, then your expected return is $½ \times 100 + ½ \times 200 = 150$. If you receive a medical treatment where:

- with probability ¼ people live for another 40 years on average,
- with probability ¾ people live for another 4 years on average,

then on average you live for another $¼ \times 40 + ¾ \times 4 = 10 + 3 = 13$ years.

In general, if prospect x has an expectation (of money, life duration, or whatever) EV(x) and prospect y has an expectation EV(y), then receiving x with probability λ and y with probability $1-\lambda$ has expectation $\lambda EV(x) + (1-\lambda)EV(y)$. This means that expectation is linear with respect to probabilistic mixing. The result holds in particular for EU, that is, not only if the expectation concerns money or life years, but also if the expectation concerns utility. If x has expected utility EU(x), and y has expected utility EU(y), then receiving x with probability λ and y with probability $1-\lambda$ has expected utility $\lambda EU(x) + (1-\lambda)EU(y)$.

Exercise 2.6.6.[a] Show that EU is *linear in probability*; i.e., $EU(x_\lambda y) = \lambda EU(x) + (1-\lambda)EU(y)$. □

We now return to the idea of McCord & de Neufville (1986). They proposed to consider equivalences such as in Figure 2.6.4, rather than the SG equivalences of Figure 2.5.3. They called the resulting measurement method the lottery-equivalent method. Basically, the same indifference is elicited as in Figure 2.5.3, but now α and the equivalent SG-prospect have been mixed with another "common" prospect C. This way they avoid the use of riskless prospects.

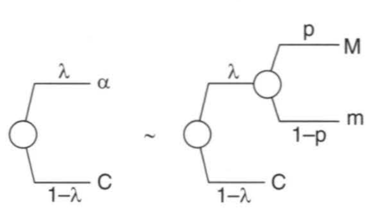

Figure 2.6.4 The lottery-equivalent method of McCord & de Neufville (1986) ($\lambda > 0$).

Because of the linearity in probability of EU (Exercise 2.6.6), the EU of the left prospect in the figure is $\lambda U(\alpha) + (1-\lambda)EU(C)$, and of the right prospect it is $\lambda(pU(M) + (1-p)U(m)) + (1-\lambda)EU(C)$. The equality of these expected utilities implies the same equality $U(\alpha) = pU(M) + (1-p)U(m) = p$ as resulting from the SG equivalence of the preceding section (the SG Eq. (2.5.1)).[5]

[5] Cancel the common term $(1-\lambda)EU(C)$, divide by the positive λ, and set again $U(m) = 0$ and $U(M) = 1$.

In a perfect world where EU would hold perfectly well, there would be no inconsistencies in observations. Then it would not matter if we use Figure 2.5.3 or Figure 2.6.4 to measure p, and p should be the same in both figures. This identity requirement is called standard-gamble (SG) consistency. Then the indifference in Figure 2.5.3 should always imply the indifference in Figure 2.6.4. This implication is stated formally in Figure 2.6.5.

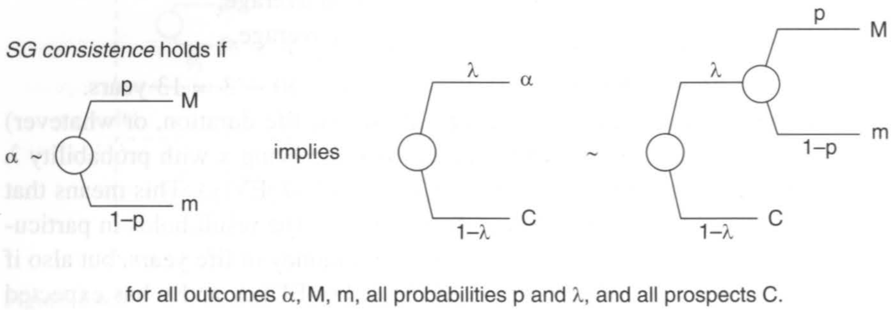

Figure 2.6.5

The condition has normative appeal. If the left indifference in Figure 2.6.5 – i.e., the indifference in Figure 2.5.3 – holds, then in the right indifference (Figure 2.6.4) you receive the same C for both prospects if going down ($1-\lambda$), and equivalent things if going up (λ), so that you are equally well off in either case. Hence, the two compound prospects seem to be equally attractive to you. The following theorem shows that SG consistency, a weak version of the well-known independence condition (see next section) together with some natural conditions, is not only necessary, but also sufficient, for EU. Von Neumann & Morgenstern (1947) presented a first version of this theorem, but their result was incomplete. Several authors have subsequently contributed to variations of the result (Fishburn 1970 p. 103). The following theorem is yet another variation.

Theorem 2.6.3 [EU for risk]. Under Structural Assumption 2.5.2 (decision under risk and richness), the following two statements are equivalent.

(i) Expected utility holds.
(ii) $\succcurlyeq$ satisfies:
- weak ordering;
- SG solvability;
- SG dominance;
- SG consistency. □

We saw, in Exercise 2.6.4, that the utility U can be replaced by any $U^* = \tau + \sigma U$ for real τ and positive σ. Here σ is the *unit parameter* because it determines the unit of measurement and τ is the *level parameter* because it affects absolute levels of utility. It turns out that no other utility functions are possible than the ones just described.

In general, a function in a particular model is *unique up to unit and level* if it can be replaced by another function if and only if the other function differs by unit and level. We also call such functions *interval scales*. The term cardinal has also been used to refer to this mathematical property, but this term has many interpretational connotations. We therefore will not use it.

Observation 2.6.3′ [Uniqueness result for Theorem 2.6.3]. In EU, utility is unique up to level and unit. □

Comment 2.6.4 [Organization of measurements and behavioral foundations; the five-step presentation]. The line of presentation leading to Theorem 2.6.3 is typical of the presentation of decision models throughout this book. It involved the following five steps: (1) The model, EU, was defined in Definition 2.5.3, with utility as subjective parameter; (2) Exercises 2.5.1 and 2.5.2 showed how to derive decisions from the model's parameters; (3) Exercises 2.3–2.5 and Eq. (2.5.1) presented the elicitation method for deriving the model's parameters from decisions; (4) §2.6 presented consistency conditions for measurements; (5) the consistency conditions for measurements gave a behavioral foundation in Theorem 2.6.3. □

Comment 2.6.5 [Variance of utility]. It may be surprising that only the probability-weighted average of utility plays a role under EU, and no other aspects of the utility distribution. Receiving a utility of 0.6 for sure is equated with receiving utility 1 with probability 0.6 and utility 0 with probability 0.4. It may seem that the risk and, for instance, the variance of the utility received should also be relevant and that they are inappropriately ignored here (Gandjour 2008 p. 1210; Rode *et al.* 1999 p. 271). The key to understanding why this is not so, and to understanding the essence of utility, lies in the SG Figure 2.5.3. Because utility is measured in this way, essentially in terms of probabilities, it becomes commensurable with probability and risk. The critical condition of EU, SG consistency, implies that the exchange rate between value and risk as measured in the SG equivalence is universally valid, also within complex decision trees. In this manner, all of probability and risk that is relevant for the preference value of prospects is captured by utility and its average. □

2.7 Independence and other preference conditions for expected utility

In the literature, many variations have been used of the preference conditions presented in the preceding section. Somewhat stronger (more restrictive) than SG consistency is the *substitution principle*:

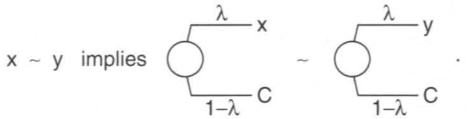

It entails that replacing a prospect in a mixture by another equally good prospect does not affect the preference value of the mixture. SG consistency concerned the special case where x was a sure outcome and y a two-outcome prospect. Yet stronger is the *independence condition*, a condition most often used in the literature for behavioral foundations of EU:

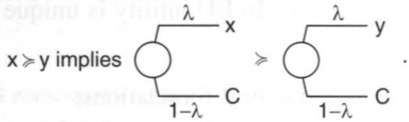

It entails that improving a prospect in a mixture improves the mixture. Substitution follows by first applying independence with x ≽ y and then with x ≼ y. A violation of independence suggests a violation of a conservation-of-preference principle, alluded to before when discussing Dutch books. Although I will not try to formalize this general idea, the following reasoning may be suggestive: If the preference between the mixtures were reversed so that replacing x by y in the mixture would lead to a strictly preferred prospect, then we would have created positive preference value by keeping C fixed and worsening x. It signals that there must be some positive interaction between C and y (or negative between C and x). Given that in no possible physical world C and x, or C and y, can coexist (you never receive both C and x, or both C and y), the interaction must be purely psychological, occurring only in the mind of the decision maker. The decision maker is, magically, creating value out of nothing but own imagination.

Independence, and its weakenings presented here, may seem to be completely self-evident at first sight, and I also think they *are* normatively compelling. However, these conditions have many implications, and many people, upon examining the implications, come to reject independence as a universal normative principle. As the psychologist Duncan Luce once put it (Luce 2000 §2.4.2), "Although this line of rational argument seems fairly compelling in the abstract, it loses its force in some concrete situations." We will return to criticisms of independence in later sections.

Because SG consistency is weaker than independence, the implication of (ii) ⇒ (i) in Theorem 2.6.3 (EU for risk) with SG consistency is a stronger result than it would be if independence were to replace SG consistency. The result with independence follows as an immediate corollary of the result with SG consistency and this does not hold the other way around. In this sense, SG consistency is a more useful preference condition.

In many applications, the implication (i) ⇒ (ii) in Theorem 2.6.3 is useful. Strong preference conditions may be interesting in such applications. For example, if we want to falsify EU empirically, then it will be easier to obtain violations of independence than of SG consistency. Because independence is stronger than SG consistency, violations of the former are easier to detect than violations of the latter. For example, if there are no indifferences in our data then SG consistency cannot be falsified in any direct manner, but independence can be. This will happen, for instance, when we discuss the Allais paradox in §4.12. This paradox directly falsifies independence, showing that EU is violated, but it does not directly falsify SG consistency. Hence, it is useful to know various preference conditions for decision theories, and not only the weakest ones that give the most general behavioral foundations.

In the presence of the other conditions for EU, SG dominance implies a very natural condition, usually considered a hallmark of rationality, namely *(first-order) stochastic dominance*: shifting probability mass from an outcome to a preferred outcome should lead to a preferred prospect. In other words, the more money the better. This preference condition extends SG dominance to general prospects with possibly more than two outcomes, using weak (unqualified) rather than strict preference. As will be explained later, the condition defined here is equivalent to other formulations used in the literature, stated in terms of distribution functions or improvements of outcomes.

Exercise 2.7.1.[!a] Show that EU implies stochastic dominance. □

Exercise 2.7.2.[b] Use Exercise 2.6.6 to show that EU implies independence. □
The following variation of independence is similar to Savage's (1954) sure-thing principle, requiring that preference be independent of common outcomes, and defined later for uncertainty (Eq. (4.8.1)). We call it the *sure-thing principle for risk*, and it is defined in Figure 2.7.1. It can be seen that independence and this condition are, in general, logically independent. In the presence of natural conditions plus a probabilistic continuity condition they become equivalent. Hence, the sure-thing principle for risk could also be used to obtain a preference foundation for EU. We do not elaborate on this point.

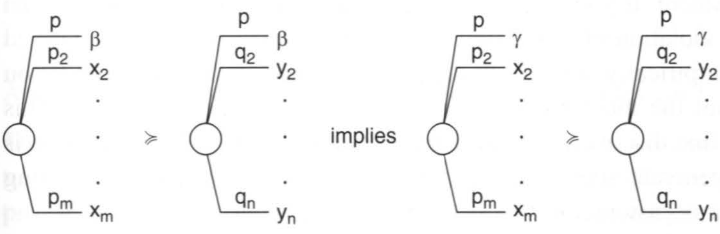

Figure 2.7.1 The sure-thing principle for risk.

Exercise 2.7.3.[!a] Show that EU implies the sure-thing principle for risk. □

2.8 Basic choice inconsistencies

The following exercise will be discussed immediately in the text that follows.

Exercise 2.8.1.[a] Substitute your own certainty equivalent for both prospects depicted in Figure 2.8.1.

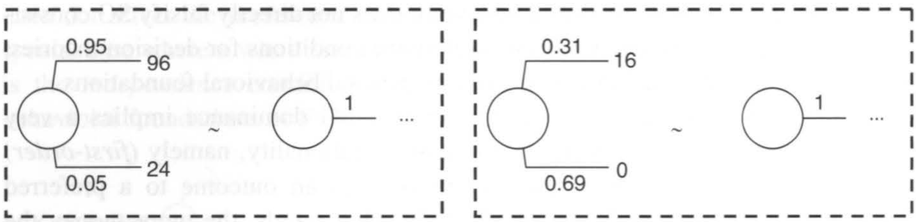

Figure 2.8.1

□

Violations of EU will be discussed in later chapters, and the choices in Figure 2.4.1 will then play a role. This section only discusses some violations of general choice principles, revealed by the choices in Figures 2.4.1, 2.4.2, and 2.8.1. The choice in Fig. 2.4.1d can be related to Figs. 2.4.2b and 2.8.1b, where the two prospects reappear. It is natural that the prospect that you preferred in Fig. 2.4.1d had the highest certainty equivalent. If this is not the case for you, then a violation of stochastic dominance or transitivity can be derived from your choices. Empirically, a majority of subjects exhibit such a violation. That is, they choose the upper prospect in Fig. 2.4.1d but assign a lower CE to it in Fig. 2.4.2b relative to Fig. 2.8.1b. These violations are known as preference reversals (Bardsley *et al.* 2010 Box 3.4).

The certainty equivalents in Figs. 2.4.2a and Fig. 2.8.1a can also be related to each other. The prospect in Fig. 2.8.1a dominates the one in Fig. 2.4.2a because the 0 outcome has been replaced by the better outcome 24. Hence, the prospect in Fig. 2.8.1a should be preferred to the one in Fig. 2.4.2a. It should, therefore, have a higher certainty equivalent. If your certainty equivalent in Fig. 2.8.1a was lower than in Fig. 2.4.2a, then a violation of stochastic dominance or transitivity can be derived from your choices. Empirically, a majority of subjects exhibit such a violation. (You were exposed to about the most tricky stimuli that this field has developed.) This choice anomaly was first discovered by Birnbaum *et al.* (1992). Their explanation is that zero outcomes generate special biases, with people sometimes overweighting their role but on other occasions, such as here, improperly ignoring them.

Empirical findings such as those just discussed are notorious because they reveal many violations of rationality requirements, such as transitivity or stochastic dominance. The prevalence of such anomalies has led some authors to conclude that studying individual preference is a hopeless task because, as they argue, such anomalies clearly show that we are not measuring "true values" in any sense, and we are only measuring random choices. Such views are captured in pessimistic versions of the constructive view of preference, discussed by Slovic (1995). There remain, fortunately, enough regularities in data to make the empirical study of decision theory worthwhile. In addition, for many decisions, individual preferences and utilities are the crucial and decisive factors. We then have to measure them as well as we can, no matter what the difficulties are. Optimistic versions of the constructive view of preference argue that we, decision analysts, aware of the deficiencies of our measurement instruments, should work and interact more with clients, "not as archaeologists, carefully uncovering what is there, but as architects, working to build a defensible expression of value" (Gregory, Lichtenstein, & Slovic 1993 p. 179). Using interactive sessions we can get more out of fewer subjects.

Whereas it is common jargon in the field to refer to choice anomalies as irrationalities and biases on the part of the decision makers, it may be argued that these anomalies are a problem of our measurement instruments rather than of the decision makers. The purported biases and irrationalities are usually discovered in experimental measurements of preference values. Such measurements, especially if based on hypothetical choice – as done in Figures 2.4.1, 2.4.2, and 2.8.1 – are hard to relate to for subjects. When answering the questions concerning these figures, you may have felt that your answers were at least in part random. Thus our measurement instruments, even if the best conceivable, do not perfectly tap into the value systems of subjects. In measurements of preference and value it is, therefore, desirable to clarify the nature of stimuli to subjects as much as possible, and to use real incentives in descriptive studies whenever possible.

For prescriptive purposes, a big value of decision theory lies in uncovering anomalies such as those just described. Such anomalies do exist in individual choices and indeed cannot reflect true values. By revealing such anomalies, decision theory can show people that in some situations they are not following their true value system, and that they can and should improve their decisions in such situations. In other words, in such situations decision theory can bring new insights and improve decisions. It can do this by using simple consistency checks, without any need to know the exact true values of a client. As pointed out by Raiffa (1961), it is not problematic but, rather, essential for prescriptive applications of decision theory that violations of the rational models exist in natural behavior.

Appendix 2.9 Proof of Theorem 2.6.3 and Observation 2.6.3′

The proof of Theorem 2.6.3 is simple. It makes a crucial achievement of EU transparent, namely that EU provides commensurability between outcomes and probability. SG measurements settle the exchange rate between outcomes and probabilities, and this exchange rate should subsequently apply to all situations. This is the idea underlying the proof.

It was explained in the main text that the preference conditions are necessary for EU. We, henceforth, assume the preference conditions and derive EU. Take two outcomes $M > m$. Define $U(m) = 0$ and $U(M) = 1$ for now. For each $m < \alpha < M$ define $U(\alpha)$ as the SG probability of α (Figure 2.5.3), which can be done because of SG solvability. Consider a prospect with all outcomes between m and M. Figure 2.9.1 is explained next.

The first equivalence follows from SG consistency, by replacing the riskless prospect x_j by its equivalent standard gamble. The second equivalence follows the same way, by repeated application. The equality is by multiplication of probabilities. Every prospect is thus equivalent to the prospect that with a probability equal to the EU of the original prospect yields outcome M, and outcome m otherwise. By SG

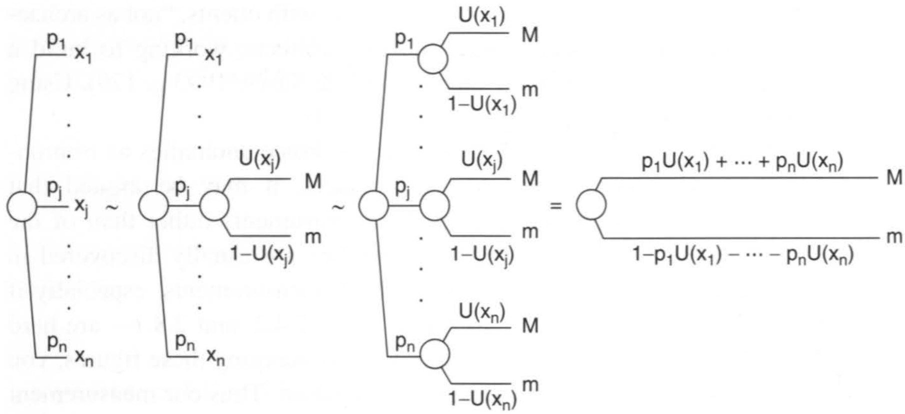

Figure 2.9.1

dominance, the latter prospects, having only m and M as outcomes, are ordered by their EU. By transitivity, so are all prospects with outcomes between m and M. We have proved that EU holds on the domain of those prospects.

For the uniqueness result in Observation 2.6.3′ applied to the domain of all prospects with outcomes between m and M, we can freely choose unit and level of U by Exercise 2.6.4. This is the only freedom we have. For any alternative U*, we can substitute $U^*(m) = \tau$ and $U^*(M) = \sigma + \tau$ for any real τ and $\sigma > 0$, after which U* is uniquely determined from SG equivalences as $\tau + \sigma U$ for the utility U with $U(m) = 0$ and $U(M) = 1$.

We can carry out the above reasoning for any $m < M$, thus obtaining EU on every bounded part of our domain and, finally, obtain EU on the whole domain by routine extensions. For the extensions of the domain (reducing m and enlarging M) it is easiest to rescale all newly defined models such that utility is 0 at the original m we started with and it is 1 at the original M we started with. Then these models all coincide on common domains by the uniqueness result just established, and they are genuine extensions of each other. Observation 2.6.3′ thus holds on the domain limited by any m and M and, consequently, on the whole domain.

By SG dominance with $p = 1$ and $q = 0$, $\alpha > \beta$ implies $\alpha \succ \beta$. This implies that U increases strictly. The proof is now complete.

The proof presented here adapts some appealing substitution-based proofs in the literature (Arrow 1951a pp. 424–425; Luce & Raiffa 1957 pp. 27–28) to our preference conditions. I constructed the proof of the bookmaking theorem in Appendix 1.11 by using the proof of this appendix and then applying a duality between states and outcomes (Abdellaoui & Wakker 2005 end of §1). Figure 1.11.1 is the analog of Figure 2.9.1. In the former figure we replaced every event E by its certainty equivalent P(E) and then used linearity in outcomes. In the latter figure we replaced every outcome by its standard-gamble probability (of getting M versus m) and then used linearity in probability. □