5 Utility

It is hard to think of any minimally reasonable decision rule that totally ignores how much utility an act brings about. However, the concept of utility has many different technical meanings, and it is important to keep these different meanings apart. In Chapter 2 we distinguished three fundamentally different kinds of measurement scales. All scales are numerical, i.e. utility is represented by real numbers, but the information conveyed by the numbers depends on which type of scale is being used.

1. Ordinal scales ("10 is better than 5")
2. Interval scales ("the difference between 10 and 5 equals that between 5 and 0")
3. Ratio scales ("10 is twice as valuable as 5")

In ordinal scales, better objects are assigned higher numbers. However, the numbers do not reflect any information about differences or ratios between objects. If we wish to be entitled to say, for instance, that the difference between ten and five units is exactly as great as the difference between five and zero units, then utility has to be measured on an interval scale. Furthermore, to be entitled to say that ten units of utility is twice as much as five, utility must be measured on a ratio scale.

Arguably, utility cannot be directly revealed by introspection. We could of course ask people to estimate their utilities, but answers gathered by this method would most certainly be arbitrary. Some more sophisticated method is needed. So how on earth is it possible to assign precise numbers to outcomes and acts that accurately reflect their value? And how can this process be a scientific one, rather than a process that is merely arbitrary and at best based on educated guesses?

5.1 How to Construct an Ordinal Scale

Let us first show how to construct an ordinal utility scale. To make the discussion concrete, imagine that you have a collection of crime novels, and that you wish to assign ordinal utilities to each book in your collection. To avoid a number of purely mathematical obstacles we assume that the number of books in the collection is finite. Now, it would by no means be unreasonable to ask you to answer the following question: "Which novel do you like the most, *The Hound of the Baskervilles* by Arthur Conan Doyle or *Death on the Nile* by Agatha Christie?" Suppose you answer *Death on the Nile*. This preference can be represented by the symbol $\succ$ in the following way.

(1) *Death on the Nile* $\succ$ *The Hound of the Baskervilles*

Proposition (1) is true if and only if *Death on the Nile* is preferred over *The Hound of the Baskervilles*. This may sound trivial, but what does this mean, more precisely? How do we check if you do in fact prefer *Death on the Nile* to *The Hound of the Baskervilles*? Economists argue that there is an easy answer to this question: Your preferences are *revealed* in your choice behavior. Therefore, you prefer x to y if and only if you choose x over y whenever given the opportunity. The main advantage of this proposal is that it links preference to directly observable behavior, which entails that the concept of preference (and hence utility) becomes firmly connected with empirical observations. However, it is of course easy to question this alleged connection between choice and preference. Perhaps you actually preferred x over y, but chose y by mistake, or did not know that y was available. Furthermore, using the behaviorist interpretation of preferences it becomes difficult to distinguish between strict preference ("strictly better than") and indifference ("equally good as"). The observation that you repeatedly choose x over y is equally compatible with the hypothesis that you strictly prefer x over y as with the hypothesis that you are indifferent between the two. However, for the sake of the argument we assume that the decision maker is able, in one way or another, to correctly state pairwise preferences between any pair of objects.

Indifference will be represented by the symbol $\sim$. For future reference we also introduce the symbol $\succeq$, which represents the relation "at least as preferred as." We now have three different preference relations, $\succ$, $\sim$ and $\succeq$. Each of them can easily be defined in terms of the others.

(2a) $x \succeq y$ if and only if $x \succ y$ or $x \sim y$

(2b) $x \sim y$ if and only if $x \succeq y$ and $y \succeq x$

(2c) $x \succ y$ if and only if $x \succeq y$ and not $x \sim y$

Let us return to the example with crime novels. The entire collection of novels can be thought of as a set $B = \{x, y, z, \ldots\}$, where book x is *Death on the Nile* and book y is *The Hound of the Baskervilles*, and so on. Since you were able to state a preference between *Death on the Nile* and *The Hound of the Baskervilles* it seems reasonable to expect that you are able to compare any two books in your collection. That is, for every x and y in B, it should hold that:

Completeness *$x \succ y$ or $x \sim y$ or $y \succ x$.*

Completeness rules out the possibility that you fail to muster any preference between some pairs of books. This may sound trivial, but it is in fact a rather strong assumption. To see this, consider a completely different kind of choice: Do you prefer to be satisfied but stupid like a pig, or dissatisfied but clever as Socrates? A natural reaction is to say that the two alternatives are incomparable. (J.S. Mill famously had a different opinion; he preferred the latter to the former.) The possibility of incomparability is not consistent with the completeness axiom.

We moreover assume that strict preferences are *asymmetric* and *negatively transitive*.

Asymmetry *If $x \succ y$, then it is false that $y \succ x$.*

Transitivity *If $x \succ y$ and $y \succ z$, then $x \succ z$.*

Negative transitivity *If it is false that $x \succ y$ and false that $y \succ z$, then it is false that $x \succ z$.*

From asymmetry we can conclude that if *Death on the Nile* is preferred over *The Hound of the Baskervilles*, then it is not the case that *The Hound of the Baskervilles* is preferred over *Death on the Nile*. Transitivity is a slightly more complex property: if *Death on the Nile* is preferred over *The Hound of the Baskervilles* and *The Hound of the Baskervilles* is preferred over *Sparkling Cyanide*, then *Death on the Nile* is preferred over *Sparkling Cyanide*. This assumption has been seriously questioned in the literature, for reasons we will return to in Chapter 8. Here, we shall accept it without further ado.

Negative transitivity is a slightly stronger version of transitivity. The reason why we assume negative transitivity rather than just transitivity is because we need it. Negative transitivity implies that indifference is transitive (if $x \sim y$ and $y \sim z$, then $x \sim z$), but this does not follow from the assumption that strict preference is transitive. (I leave it to the reader to verify this claim.) Furthermore, it can also be shown that negative transitivity, but not transitivity, entails that $x \succ z$ if and only if, for all y in B, $x \succ y$ or $y \succ z$.

Let us now return to the problem of constructing an ordinal utility scale. What assumptions must we make about preference relations for this to be possible? That is, what must we assume about preferences for there to exist a function u that assigns real numbers to all books in your collection such that better books are assigned higher numbers, i.e.

(3) $x \succ y$ if and only if $u(x) > u(y)$?

Note that there would not, for instance, exist a utility function u if the strict preference relation is cyclic, i.e. if $x \succ y$ and $y \succ z$ and $z \succ x$. There are simply no real numbers that can represent this ordering such that better

Box 5.1 An ordinal utility scale

Theorem 5.1 Let B be a finite set of objects. Then, $\succ$ is complete, asymmetric and negatively transitive in B if and only if there exists a real-valued function u such that condition (3) on this page.

Proof We have to prove both directions of the biconditional. Let us first show that *if* there exists a function u such that (3) holds, *then* the preference relation $\succ$ is complete, asymmetric and negatively transitive. Completeness follows directly from the corresponding property of the real numbers: Since it is always true that $u(x) > u(y)$ or $u(x) = u(y)$ or $u(y) > u(x)$, it follows that $x \succ y$ or $x \sim y$ or $y \succ x$. Furthermore, note that if $u(x) > u(y)$, then it is not the case that $u(y) > u(x)$; this is sufficient for seeing that $\succ$ is asymmetric. Finally, negative transitivity follows from the observation that if $u(x) \geq u(y)$ and $u(y) \geq u(z)$, then $u(x) \geq u(z)$. (Since $x \succ y$ is not compatible with $y \geq x$, given completeness.)

Next, we wish to show that *if* the preference relation $\succ$ is complete, asymmetric and transitive, then there exists a function u such that (3) holds. We do this by outlining one of the many possible ways in which u can be constructed. Before we start, we label the elements in B with index numbers. This can be done by just adding indices to the elements of B, i.e. $B = \{x_1, y_2, z_3 \ldots\}$. The next step is to define the set of elements that are *worse than* an arbitrarily chosen element x. Let $W(x) = \{y\colon x \succ y\}$, i.e. y is an element in W that is worse than x. Furthermore, let $N(x)$ be the set of index numbers corresponding to W, i.e. $N(x) = \{n\colon x_n \in W(x_n)\}$.

The utility function u can now be defined as follows: $u(x) = \sum_{n \in N(x)} \left(\frac{1}{2}\right)^n$

All that remains to do is to verify that u satisfies (3), i.e. that $x \succ y$ if and only if $u(x) > u(y)$. From left to right: Suppose that $x \succ y$. Then, since $\succ$ is transitive and asymmetric, $W(x)$ must contain at least one element that is not in $W(y)$. Hence, $N(x)$ must also contain at least one element that is not in $N(y)$. By stipulation we now have:

$$u(x) = \sum_{n \in N(x)} \left(\frac{1}{2}\right)^n = \sum_{n \in N(y)} \left(\frac{1}{2}\right)^n + \sum_{n \in N(x)-N(y)} \left(\frac{1}{2}\right)^n > \sum_{n \in N(y)} \left(\frac{1}{2}\right)^n = u(y)$$

From right to left: Suppose that $u(x) > u(y)$. Completeness guarantees that $x \succ y$ *or* $x \sim y$ *or* $y \succ x$. The last alternative is impossible, since it has already been shown above that if $y \succ x$ then $u(y) > u(x)$, which is a contradiction. For the same reason, we can rule out $x \sim y$, because $u(x) = u(y)$ is also inconsistent with $u(x) > u(y)$. Hence, we can conclude that $x \succ y$. □

objects are represented by higher numbers. However, it can be shown that such numbers do exist, i.e. that there is a real-valued function u such that (3) holds true, if and only if the relation $\succ$ is complete, asymmetric and negatively transitive. (A proof is given in Box 5.1.) This means that we have solved the problem we set out to solve – we now know under what conditions it is possible to construct an ordinal utility scale, and how to do it!

5.2 Von Neumann and Morgenstern's Interval Scale

In many cases ordinal utility scales do not provide the information required for analyzing a decision problem. The expected utility principle as well as several other decision rules presuppose that utility is measured on an interval scale. In the second edition of the *Theory of Games and Economic Behavior*, published in 1947, John von Neumann and Oskar Morgenstern proposed a theory that has become the default strategy for constructing interval scales, with which other theories of utility must compete. This is partly because von Neumann and Morgenstern's theory does not merely explain what utility is and how it can be measured; their theory also offers an indirect but elegant justification of the principle of maximizing expected utility.

The key idea in von Neumann and Morgenstern's theory is to ask the decision maker to state a set of preferences over risky acts. These acts are called lotteries, because the outcome of each act is assumed to be randomly determined by events (with known probabilities) that cannot be controlled by the decision maker. The set of preferences over lotteries is then used for calculating utilities by reasoning "backwards." To illustrate this point, we will consider a simple example. Note that the example only gives a partial and preliminary illustration of von Neumann and Morgenstern's theory.

Suppose that Mr. Simpson has decided to go to a rock concert, and that there are three bands playing in Springfield tonight. Mr. Simpson thinks that band A is better than band B, which is better than band C. For some reason, never mind why, it is not possible to get a ticket that allows him to watch band A or C with 100% certainty. However, he is offered a ticket that gives him a 70% chance of watching band A and a 30% chance of watching band C. The only other ticket available allows him to watch band B with 100% certainty. For simplicity, we shall refer to both options as lotteries, even though only the first involves a genuine element of chance. After considering his two options carefully, Mr. Simpson declares them to be equally attractive. That is, to watch band B with 100% certainty is, for Mr. Simpson, exactly as valuable as a 70% chance of watching band A and a 30% chance of watching band C. Now suppose we happen to know that Mr. Simpson always acts in accordance with the principle of maximizing expected utility. By reasoning backwards, we can then figure out what his utility for rock bands is, i.e. we can determine the values of $u(A)$, $u(B)$ and $u(C)$. Consider the following equation, which formalizes the hypothesis that Mr. Simpson accepts the expected utility principle and thinks that the two options are equally desirable:

$$0.7 \cdot u(\text{A}) + 0.3 \cdot u(C) = 1.0 \cdot u(\text{B}) \tag{1}$$

Equation (1) has three unknown variables, so it has infinitely many solutions. However, if utility is measured on an interval scale, Equation (1) nevertheless provides all the information we need, since the unit and end points of an interval scale can be chosen arbitrarily. (See Chapter 2.) We may therefore stipulate that the utility of the best outcome is 100 (that is, $u(\text{A}) = 100$), and that the utility of the worst outcome is 0 (that is, $u(\text{C}) = 0$). By inserting these arbitrarily chosen end points into Equation (1), we get the following equation.

$$0.7 \cdot 100 + 0.3 \cdot 0 = 1.0 \cdot u(\text{B}) \tag{2}$$

Equation (2) has only one unknown variable, and it can be easily solved: As you can see, $u(\text{B}) = 70$. Hence, we now know the following.

$$\begin{aligned} u(A) &= 100 \\ u(B) &= 70 \\ u(C) &= 0 \end{aligned}$$

Now suppose that Mr. Simpson is told that another rock band will play in Springfield tonight, namely band D, which he thinks is slightly better than B. What would Mr. Simpson's numerical utility of watching band D be? For some reason, never mind why, Mr. Simpson is offered a ticket that entitles him to watch band D with probability p and band C with probability $1 - p$, where p is a variable that he is free to fix himself. To figure out what his utility for D is, he asks himself the following question: "Which value of p would make me feel totally indifferent between watching band B with 100% certainty, and D with probability p and C with probability $1 - p$?" After considering his preferences carefully, Mr. Simpson finds that he is indifferent between a 100% chance of watching band B and a 78% chance of watching D combined with a 22% chance of watching C. Since we know that $u(\text{B}) = 70$ and $u(\text{C}) = 0$, it follows that:

$$1.0 \cdot 70 = 0.78 \cdot u(\text{D}) + 0.22 \cdot 0 \tag{3}$$

By solving this equation, we find that $u\ (\text{D}) = 70/0.78 \approx 89.7$. Of course, the same method could be applied for determining the utility of every type of good. For instance, if Mr. Simpson is indifferent between a 100% chance of watching rock band B and a 95% chance of winning a holiday in Malibu combined with a 5% chance of watching band C, then his utility of a holiday in Malibu is $70/0.95 = 73.7$.

An obvious problem with this preliminary version of von Neumann and Morgenstern's theory is that it presupposes that the decision maker chooses in accordance with the principle of maximizing expected utility. We seem to have no reason for thinking that the decision maker will apply the expected utility principle, rather than some other principle, for evaluating lotteries. This very strong assumption needs to be justified in one way or another. Von Neumann and Morgenstern proposed a clever way of doing that. Instead of directly assuming that the decision maker will always apply the expected utility principle (as we did above), they proposed a set of *constraints on rational preferences* which imply that the decision maker behaves *as if* he or she is making decisions by calculating expected utilities. More precisely put, von Neumann and Morgenstern were able to prove that if a decision maker's preferences over the sort of lotteries exemplified above satisfy a number of formal constraints, or *axioms*, then the decision maker's choices can be represented by a function that assigns utilities to lotteries (including lotteries comprising no uncertainty), such that one lottery is preferred to another just in case the expected utility of the first lottery exceeds that of the latter. In the remaining paragraphs of this section we shall spell out the technical assumptions underlying von Neumann and Morgenstern's theory in more detail.

We assume that Z is a finite set of basic prizes, which may include a holiday in Malibu, a ticket to a rock concert, as well as almost any kind of good. That is, the elements of Z are the kind of things that typically constitute outcomes of risky decisions. We furthermore assume that L is the set of lotteries that can be constructed from Z by applying the following inductive definition. (Note that even a 100% chance of winning a basic prize counts as a "lottery" in this theory.)

1. Every basic prize in Z is a lottery.
2. If A and B are lotteries, then so is the prospect of getting A with probability p and B with probability $1 - p$, for every $0 \leq p \leq 1$.
3. Nothing else is a lottery.

For simplicity, the formula ApB will be used as an abbreviation for a lottery in which one wins A with probability p and B with probability

$1 - p$. Thus, the second condition stated above could equally well be formulated as follows: If A and B are lotteries, then so is ApB, for every $0 \leq p \leq 1$. Furthermore, since ApB is a lottery, it follows that also $Cq(ApB)$, $0 \leq q \leq 1$ is a lottery, given that q is a probability and C is a lottery. And so on and so forth.

The next assumption introduced by von Neumann and Morgenstern holds that the decision maker is able to state pairwise preferences between lotteries. The formula $A \succ B$ means that lottery A is preferred over lottery B, and $A \sim B$ means that lottery A and B are equally preferred. Now, it should be obvious that preferences have to satisfy *some* structural conditions. For example, it would make little sense to prefer A to B *and* B to A; that is, for a rational decision maker it cannot hold true that $A \succ B$ and $B \succ A$. This property of rational preferences is usually referred to as the asymmetry condition. What further conditions could we impose upon rational preferences? Von Neumann and Morgenstern, as well as many other decision theorists, take completeness and transitivity to be two intuitively reasonable conditions. We recognize them from the discussion of ordinal utility. However, note that the objects over which one is supposed to state preferences are now a set of lotteries, not a set of certain outcomes.

vNM 1 (Completeness) *$A \succ B$ or $A \sim B$ or $B \succ A$*

vNM 2 (Transitivity) *If $A \succ B$ and $B \succ C$, then $A \succ C$*

To state the next axiom, let p be some probability strictly greater than zero.

vNM 3 (Independence) *$A \succ B$ if and only if $ApC \succ BpC$*

The independence axiom is best illustrated by considering an example. Imagine that you are offered a choice between lotteries A and B in Table 5.1. Each ticket is equally likely to be drawn, so the probability of winning, say, \$5 million if lottery B is chosen is 10/11.

Suppose that you prefer lottery A to lottery B. Then, according to the independence axiom it must also hold that you prefer the first lottery to the second in the situation illustrated in Table 5.2. That is, you must prefer ApC to BpC.

Table 5.1

	Ticket no. 1	Ticket no. 2–11
A	$1 m	$1 m
B	$0	$5 m

Table 5.2

	Ticket no. 1	Ticket no. 2–11	Ticket no. 12–100
ApC	$1 m	$1 m	$1 m
BpC	$0	$5 m	$1 m

Table 5.3

	Ticket no. 1	Ticket no. 2–11	Ticket no. 12–100
ApC	$1 m	$1 m	$0
BpC	$0	$5 m	$0

The strongest objection to the independence axiom is that it entails the Allais paradox, discussed in Chapter 4. To see this, note that the independence axiom is supposed to hold no matter what *C* is. Thus, if you prefer *A* to *B* in the first situation you must also prefer *ApC* to *BpC* in the situation illustrated in Table 5.3. However, in this case it seems entirely reasonable to hold the opposite preference, that is, to prefer *BpC* to *ApC*. This is because there is no longer any alternative that gives you a million for sure. Hence, it might be worth taking a slightly larger risk and hope to get $5 million instead.

However, despite the worries raised above we shall nevertheless suppose that the independence axiom can be defended in some way or another. This is because we need it for formulating the utility theory proposed by von Neumann and Morgenstern.

The fourth and last axiom proposed by von Neumann and Morgenstern is a continuity condition. Let p and q be some probabilities strictly greater than 0 and strictly smaller than 1.

vNM 4 (Continuity) *If* $A \succ B \succ C$ *then there exist some* p *and* q *such that* $ApC \succ B \succ AqC$

The following example explains the assumption articulated by the continuity axiom. Suppose that A is a prize worth \$10 m, B a prize worth \$9 m and C a prize worth \$0. Now, according to the continuity axiom, if you prefer \$10 m to \$9 m and \$9 m to \$0, then there must be some probability p, which may be very close to 1, such that you prefer \$10 m with probability p and \$0 with probability $1 - p$ over \$9 m for sure. Furthermore, there must be some probability q such that you prefer \$9 m for sure over \$10 m with probability q and \$0 m with probability $1 - q$. Of course, some people might feel tempted to deny that these probabilities exist; perhaps it could be argued that there is no probability p simply because \$9 m for sure is always better than a lottery yielding either \$10 m or \$0, no matter how small the probability of getting \$0 is. The standard reply to this complaint is that p might lie very close to 1.

In addition to the four axioms stated above, we also need to make an additional technical assumption, saying that the probability calculus applies to lotteries. (In some presentations this condition is listed as a separate axiom.) The essence of this assumption is that it does not matter if you are awarded prize A if you first roll a die and then roll it again, or make a double roll, provided that you only get the prize if you get two sixes. Put into mathematical vocabulary, compound lotteries can always be reduced to simple ones, involving only basic prizes. Hence, if p, q, r and s are probabilities such that $pq + (1 - p)r = s$, then $(AqB)p(ArB) \sim AsB$.

The axioms stated above imply the following theorem, which is frequently referred to as von Neumann and Morgenstern's theorem. It consists of two parts, a *representation* part and a *uniqueness* part.

Theorem 5.2 The preference relation $\succ$ satisfies vNM 1–4 if and only if there exists a function u that takes a lottery as its argument and returns a real number between 0 and 1, which has the following properties:

(1) $A \succ B$ if and only if $u(A) > u(B)$.
(2) $u(ApB) = pu(A) + (1 - p)u(B)$.
(3) For every other function u' satisfying (1) and (2), there are numbers $c > 0$ and d such that $u' = c \cdot u + d$.

Property (1) articulates the fact that the utility function u assigns higher utility numbers to better lotteries. From (1) it follows that $A \sim B$ if and only if $u(A) = u(B)$. Here is a proof: To prove the implication from left to right, suppose for reductio that $A \sim B$ and that $u(A) > u(B)$. It then follows from (1) that $A \succ B$, which contradicts the completeness axiom. Moreover, if $A \sim B$ and $u(B) > u(A)$ it follows that $B \succ A$, which also contradicts the completeness axiom. Hence, if $A \sim B$ then it has to be the case that $u(A) = u(B)$. Furthermore, to prove the implication from right to left, suppose that $u(A) = u(B)$ and that it is false that $A \sim B$. The completeness axiom then entails that either $A \succ B$ or $B \succ A$ (but not both), and in conjunction with (1) both possibilities give rise to a contradiction, because neither $u(A) > u(B)$ nor $u(B) > u(A)$ is consistent with the assumption that $u(A) = u(B)$.

Property (2) of the theorem is the expected utility property. It shows us that the value of a compound lottery equals the expected utility of its components. This means that anyone who obeys the four axioms acts *in accordance with* the principle of maximizing expected utility. Of course, it does not follow that the decision maker *consciously applies* the expected utility principle. All that follows is that it is possible to rationalize the decision maker's behavior by pointing out that he acts *as if* he or she were ascribing utilities to outcomes and calculating the expected utilities.

Properties (1) and (2) are the representation part of the theorem, which show how a utility function can be used for *representing* the decision maker's behavior. Property (3) is the uniqueness part, telling us that all utility functions satisfying (1) and (2) have something important in common: they are all positive linear transformations of each other. This means that every utility function satisfying (1) and (2) can be obtained from every other such function by multiplying the latter by a constant and adding another constant. As explained in Chapter 2, this means that in von Neumann and Morgenstern's theory utility is measured on an interval scale. A proof of von Neumann and Morgenstern's theorem can be downloaded from www.martinpeterson.org.

Commentators have outlined at least three general objections to von Neumann and Morgenstern's result. If correct, these objections show that a utility function cannot be constructed in the way proposed by them.

(1) *The axioms are too strong.* As pointed out above, the axioms on which the theory relies are not self-evidently true. It can be questioned whether

rational agents really have to obey these axioms. We shall return to this criticism, which is by far the most common one, in Chapter 8.

(2) *No action guidance.* By definition, a rational decision maker who is about to choose among a large number of very complex acts (lotteries) has to know already from the beginning which risky act (lottery) to prefer. This follows from the completeness axiom. Hence, a utility function derived in von Neumann and Morgenstern's theory cannot be used by the decision maker for first *calculating* expected utilities and thereafter *choosing* an act having the highest expected utility. The output of von Neumann and Morgenstern's theory is not a set of preferences over acts – on the contrary, these preferences are used as input to the theory. Instead, the output is a (set of) utility function(s) that can be used for describing the agent as an expected utility maximizer. Hence, ideal agents do not prefer an act *because* its expected utility is favorable, but can only be described *as if* they were acting from this principle. To some extent, the theory thus puts the cart before the horse.

In reply to this objection, it could be objected that someone who is not fully rational (and thus does not have a complete preference ordering over lotteries) might nevertheless get some help from the axioms. First, they can be used for detecting any inconsistencies in the decision maker's set of preferences. Second, once your utility function has been established the expected utility principle could be used for filling any "missing gaps," i.e. lotteries you have not yet formed preferences about. Note that both these responses presuppose that the decision maker is a nonideal agent. But what about ideal decision makers? Does it really make sense to *define* an ideal decision maker such that it becomes trivially true that ideal decision makers do not need any action guidance?

(3) *Utility without chance.* It seems rather odd from a linguistic point of view to say that the *meaning* of utility has something to do with preferences over lotteries. For even a decision maker who (falsely) believes that he lives in a world in which every act is certain to result in a known outcome, i.e. a world that is fully deterministic and known to be so, can meaningfully say that the utility of some events exceeds that of others. In everyday contexts the concept of utility has no conceptual link to the concept of risk. Hence, it might be questionable to develop a technical notion of

utility that presupposes such a link, at least if it is meant to be applied in normative contexts. (Perhaps it might be fruitful for descriptive purposes.)

The obvious reply to this objection is that von Neumann and Morgenstern's theory is not a claim about the meaning of utility, it is a claim about how to measure utility. However, this is not a very helpful reply, because it then follows that their theory is at best a partial theory. If we acknowledge that it would make sense to talk about utilities even if the world was fully deterministic and known to be so, it seems to follow that we would then have to come up with some other method for measuring utility in that world. And if such a method exists, why not use it everywhere?

Box 5.2 How to buy a car from a friend without bargaining

Joanna has decided to buy a used car. Her best friend Sue has a yellow Saab Aero Convertible that she is willing to sell to Joanna. However, since Joanna and Sue are good friends and do not want to jeopardize their friendship, they feel it would be unethical to bargain about the price. Instead, they agree that a fair price of the Saab would be the price at which Joanna is *indifferent* between the Saab and the amount of money paid for the car, irrespective of Sue's preference. Hence, if Joanna is indifferent between the Saab and $10,000 it follows that $10,000 is a fair price for the car. Unfortunately, because Joanna and Sue are friends, Joanna is not able to honestly and sincerely make direct comparisons between the Saab and various amounts of money. When asked to state a preference between the car and some amount of money she cannot tell what she prefers. Therefore, to overcome Joanna's inability to directly compare the Saab with money they decide to proceed as follows:

1. Sue offers Joanna to state preferences over a large number of hypothetical car lotteries. The prizes include a Ford, the Saab and a Jaguar. It turns out that Joanna's preferences over car lotteries satisfy the von Neumann–Morgenstern axioms and that she is indifferent between getting the Saab for certain and a lottery in which the probability is 0.8 that she wins a Ford and 0.2 that she wins

a Jaguar. The von Neumann–Morgenstern theorem (Theorem 5.2) now implies that:

(i) $0.8 \cdot u(\text{Ford}) + 0.2 \cdot u(\text{Jaguar}) = 1.0 \cdot u(\text{Saab})$
By letting $u(\text{Jaguar}) = 1$ and $u(\text{Ford}) = 0$ it follows that $u(\text{Saab}) = 0.2$

2. Next, Sue helps Joanna to establish a separate utility function of money by offering her a second set of hypothetical lotteries. This utility function has no (direct) relation to her utility of cars. As before, Joanna's preferences satisfy the von Neumann–Morgenstern axioms, and for future reference we note that she is indifferent between getting \$25,000 for certain and a lottery in which the probability is 0.6 that she wins \$60,000 and 0.4 that she wins nothing. Joanna is also indifferent between getting \$60,000 for certain and a lottery in which the probability is 0.2 that she wins \$25,000 and 0.8 that she wins \$90,000. Hence, we have:

(ii) $0.6 \cdot u(\$60{,}000) + 0.4 \cdot u(\$0) = 1.0 \cdot u(\$25{,}000)$
(iii) $0.2 \cdot u(\$25{,}000) + 0.8 \cdot u(\$90{,}000) = 1.0 \cdot u(\$60{,}000)$

Equation (ii) entails that if $u(\$60{,}000) = 1$ and $u(\$0) = 0$, then $u(\$25{,}000) = 0.6$. Furthermore, according to equation (iii), if $u(\$90{,}000) = 1$ and $u(\$25{,}000) = 0$, then $u(\$60{,}000) = 0.8$. The two utility scales derived from (ii) and (iii) are not directly connected to each other. However, they can of course be merged into a single scale by observing that the difference in utility between \$60,000 and \$25,000 corresponds to 0.4 units on scale (ii) and 0.8 units on scale (iii). This means that a difference of 0.2 units on scale (iii) between \$60,000 and \$90,000 would correspond to a difference of 0.1 unit on scale (ii). Hence, $u(\$90{,}000) = 1.1$ on that scale. The leftmost columns of Table 5.4 summarize the two original scales. The scales to the right show how the original scales can be merged into a single scale as described above, and thereafter calibrated such that the end points become 1 and 0.

3. In the third and final step Joanna has to find a way of connecting her utility scale for cars with her utility scale for money.

Table 5.4

	Scale (ii)	Scale (iii)	Merged scale	Calibrated scale
\$90,000		1	1.1	1.1/1.1 = 1
\$60,000	1	0.8	1	1/1.1 = 0.91
\$25,000	0.6	0	0.6	0.6 /1.1 = 0.55
\$0	0		0	0/1.1 = 0

As mentioned above, she finds it impossible to *directly* compare the Saab with money, but she is willing to compare other cars with money. It turns out that she is indifferent between a Jaguar and \$90,000 and between a Ford and \$25,000. So on the calibrated scale a Jaguar is worth 1 unit of utility and a Ford 0.55 units. Hence, the difference between a Jaguar and a Ford is 1 - 0.55 = 0.45 on the calibrated scale. Now recall the first step, in which we established that the utility of the Saab is 20% of the difference between the Jaguar and the Ford. As 20% of 0.45 is 0.09, the utility of the Saab is 0.45 + 0.09 = 0.54.

What amount of money corresponds to 0.54 units of utility? This is equivalent to asking: What amount \$X for certain is judged by Joanna to be exactly as attractive as a lottery in which she wins \$90,000 with a probability of 0.54 and \$0 with a probability of 0.46? As pointed out in Step 2, Joanna is indeed willing to answer this type of question. After some reflection, she concludes that the amount in question is \$55,000. This answers the question we set out to answer at the beginning: An ethically fair price of Sue's yellow Saab Aero Convertible is \$55,000.

In summary, this example shows how the von Neumann–Morgenstern theory can be used for making *indirect* comparisons of items the agent is not immediately willing to compare, e.g. a specific car and some amount of money. The trick is to construct two separate scales, one for cars and one for money, and then weld them together into a single scale. So in this type of case Joanna's preferences satisfy the completeness axiom in an *indirect* sense.

5.3 Can Utility be Measured on a Ratio Scale?

Many decision theorists believe that utility can be measured only on an interval scale, and that the best way to construct an interval scale is along the lines suggested by von Neumann and Morgenstern. Arguably, this is a metaphysical claim about the nature of utility. If correct, it tells us something important about what utility is. However, some economists and mathematicians maintain that it is also possible to measure utility on a ratio scale. To render this claim plausible, a radically different approach has to be taken. The point of departure is the observation that there seems to be an intimate link between the utility of an option and the probability with which it is chosen. At first glance this may seem trivial. If you think that \$20 is better than \$10, then you will of course choose \$20 with probability 1 and \$10 with probability 0, whenever offered a choice between the two objects. However, the *probabilistic theory* of utility, as I shall call it, is much more sophisticated than that.

To start with, let us suppose that an external observer wishes to figure out what a decision maker's utility function for some set of objects is. We assume that it is not entirely sure, from the observer's perspective, what the decision maker will choose. All the observer can tell is that the probability that some option will be chosen over another is, say, 0.8. Imagine, for instance, that the manager of a supermarket observes that one of her customers tends to buy apples but no bananas eight times out of ten, and bananas but no apples two times out of ten. In that case the manager may conclude that the customer will choose an apple over a banana with a probability of 0.8. Now assume that the observer is somehow able to verify that none of the customer's tastes or desires has changed; that is, the reason why the customer sometimes buys apples and sometimes buys bananas is not that he frequently changes his taste.

Empirical studies show that probabilistic choice behavior is by no means uncommon, but how can such behavior be rational? It has been suggested that we may think of the decision maker's choice process as a two-step process. In the first step, the decision maker assesses the utility of each option. Then, in the second step, she makes a choice among the available options by simply trying to choose an alternative that maximizes utility. Now, there are at least two possibilities to consider, each of which gives rise to probabilistic choice behavior. First, the decision maker may sometimes fail to choose an option that maximizes utility in the second step of the choice process, because she fails to choose the alternative that is best for her. This version of the probabilistic theory is called the *constant utility model*. Second, we may imagine that the decision maker in the second step always chooses an alternative with the highest utility, although the first step is probabilistic, i.e. the probabilities come from the assessment of utilities rather than the choice itself. This version of the probabilistic theory is called the *random utility model*. According to this view one and the same object can be assigned different utilities at two different times, even if everything else in the world is kept constant, because the act of assigning utilities to objects is essentially stochastic.

Now, in order to establish a more precise link between probability and utility, it is helpful to formalize the probabilistic approach. Let x, y, z be arbitrary objects, and let A, B, ... be sets of objects, and let the formula $p(x \succ B)$ denote a probability of p that x will be chosen out of B. Hence, $p\ (apple \succ \{apple; banana\}) = 0.8$ means that the probability is 0.8 that an apple will be chosen over a banana. Furthermore, let $p(A \succ B)$ denote the probability of the following conditional: If A is a subset of B, then the probability is p that the chosen alternative in B is also an element of A.

We now come to the technical core of the probabilistic theory. This is the choice axiom proposed by Duncan Luce. This axiom holds that if A is a subset of B, then the probability that x will be chosen from B equals the probability that x will be chosen from A multiplied by the probability that the chosen alternative in B is also an element of A. In symbols,

Choice axiom if $A \subset B$, then $p(x \succ B) = p(x \succ A) \cdot p(A \succ B)$.

To grasp what kind of assumption is at stake here, it is helpful to consider an example. Suppose that Mr. Simpson is visiting a posh restaurant and that he is about to choose a wine from a list containing two red and two white wines. Now, the choice axiom entails that it should not matter if Mr. Simpson divides his choice into two stages, that is, first chooses between red and white wine and then between the wines in the chosen subset, or chooses directly which of the four wines to order. Hence, if Mr. Simpson is indifferent between red and white wine in general, as well as between the two red wines and the two white ones at hand, the probability that a particular bottle will be chosen is $\frac{1}{2} \cdot \frac{1}{2} = \frac{1}{4}$.

Gerard Debreu, winner of the Nobel Prize in Economics, has constructed an interesting counterexample to the choice axiom: When having dinner in Las Vegas, Mr. Simpson is indifferent between seafood and meat, as well as between steak and roast beef. The menu comprises only three dishes: x = lobster, y = steak and z = roast beef. Let B be the entire menu, let A be the set comprising x and y, and let A' be the set of y and z. Then, since Mr. Simpson is indifferent between seafood and meat, $p(x \succ B) = \frac{1}{2}$. However, $p(A \succ B) = 1 - p(z \succ B) = 1 - p(z \succ A') \cdot p\left(A' \succ B\right) = 1 - \frac{1}{2} \cdot \frac{1}{2} = \frac{3}{4}$. Hence, the axiom implies that $p(x \succ B) = p(x \succ A) \cdot p(A \succ B) = \frac{1}{2} \cdot \frac{3}{4} = \frac{3}{8}$.

In response to Debreu's example, people wishing to defend the choice axiom may say that the root of the problem lies in the individuation of alternatives. It could be argued that lobster, steak and roast beef are not alternatives at the same level. Lobster belongs to the category "seafood," and could equally well be replaced with tuna, or any other seafood dish. However, for the example to work, neither steak nor roast beef can be replaced with some other meat dish, say kebab, because then it would no longer be certain that the agent will remain indifferent between the two meat dishes. Hence, the moral of Debreu's example seems to be that alternatives must be individuated with care. Perhaps the choice axiom should be taken into account already when alternatives are being individuated – it could be conceived of as a normative requirement for how alternatives *ought* to be individuated. We are now in a position to prove the following theorem.

Theorem 5.3 Let B be a finite set of objects such that $p(x \succ B) \neq 0; 1$ for all x in B. Then, if the choice axiom holds for B and all its subsets, and the axioms of probability theory hold, then there exists a positive real-valued function u on B such that for every $A \subset B$ it holds that

(1) $p(x \succ A) = \dfrac{u(x)}{\sum_{y \in A} u(y)}$

(2) for every other function u' satisfying condition (1) there is a constant k such that $u = k \cdot u'$.

Proof First consider the existence part, (1), saying that the utility function u exists. Since it was assumed that the probability that x is chosen is nonzero, it also holds that $p(A \succ B) \neq 0$. It then follows that $p(x \succ A) = \frac{p(x \succ B)}{p(A \succ B)}$.

We stipulate that $u\ (x) = k \cdot p(x \succ B)$, where $k > 0$. Then, since the elements of A are mutually exclusive, the axioms of the probability calculus guarantee that $p(x \succ A) = \frac{k \cdot p(x \succ B)}{\sum_{y \in A} k \cdot p(y \succ B)} = \frac{u(x)}{\sum_{y \in A} u(y)}$. We now move on to the uniqueness part, (2). Suppose that u' is another utility function defined as above and satisfying (1). Then, for every x in B, it holds that $u(x) = k \cdot p(x \succ B) = \frac{k \cdot u'(x)}{\sum_{y \in A} u'(y)}$. □

The downside of the probabilistic theory is that it requires that $p \neq 0$, i.e. that each option is chosen with a nonzero probability. (This assumption is essential in the proof of Theorem 5.3, because no number can be divided by 0.) However, as pointed out above, the probability that you choose \$10 rather than \$20 if offered a choice between the two is likely to be 0. Call this the problem of perfect discrimination. The problem of perfect discrimination can be overcome by showing that for every set of objects B there exists some *incomparable* object x^* such that $p(x^* \succ x) \neq 0, 1$ for every x in B. Suppose, for example, that I wish to determine my utility of \$20, \$30 and \$40, respectively. In this case, the incomparable object can be taken to be a photo of my beloved cat Carla, who died when I was fourteen. The photo of this cat has no precise monetary value for me; my choice between money and the photo is always probabilistic. If offered a choice between \$20 and the photo, the probability is 1/4 that I would choose the money; if offered a choice between \$30 and the photo, the probability is 2/4 that I would choose the money; and if offered a choice between \$40 and the photo, the probability is 3/4 that I would choose the money. This information is sufficient for constructing a single ratio scale for all four objects. Here is how to do it.

We use the three local scales as our point of departure. They have one common element: the photo of Carla. The utility of the photo is the same in all three pairwise choices. Let $u(\text{photo}) = 1$. Then the utility of money is calculated by calibrating the three local scales such that $u(\text{photo}) = 1$ in all

Table 5.5

	u_1	u_2	u_3	u
\$20	1/4	–	–	$\frac{1/4}{3/4} = 1/3$
\$30	–	2/4	–	$\frac{2/4}{2/4} = 1$
\$40	–	–	3/4	$\frac{3/4}{1/4} = 3$
photo	3/4	2/4	1/4	1

of them. This is achieved by dividing the probability numbers listed above by 3/4, 2/4 and 1/4, respectively. Table 5.5 summarizes the example. The symbols u_1, u_2 and u_3 denote the three local scales and u denotes the single scale obtained by "welding together" the local scales.

Of course, there might exist some large amount of money that would make me choose the money over the photo with probability one. This indicates that the photo is not incomparable with every possible amount of money. However, this difficulty can be overcome by choosing some other beloved object to compare with, e.g. the only remaining photo of my daughter, or peace in the Middle East.

5.4 Can We Define Utility Without Being Able to Measure It?

Intuitively, it seems plausible to separate the meaning of the term utility from the problem of how to measure it. Consider the following analogy: We all know what it means to say that the mass of Jupiter exceeds that of Mars, but few of us are able to explain how to actually measure the mass of a planet. Therefore, if someone proposed a measurement procedure that conflicts with our intuitive understanding of mass, we would have reason to reject that procedure insofar as we wish to measure the thing we call mass. We are, under normal circumstances, not prepared to replace our intuitive concept of mass with some purely technical concept, even if the technical concept simplifies the measurement process. Does the analogous point apply to utility?

Theories of utility are sometimes interpreted as *operational definitions*, i.e. as definitions that fix the meaning of a term by setting up an empirical method for observing the entity the term is referring to. In this view, it does not make sense to distinguish meaning from measurement. As long as the concept of utility is merely used for descriptive purposes, i.e. for predicting and explaining choices, this operational approach seems fine. However, when we consider normative applications it is far from clear that an operational definition is what we are looking for. If we, for example, wish to say that a decision was rational *because* it maximized expected utility, it seems essential that the notion of utility we refer to is the true, *core notion* of utility. So what is this core notion of utility, with which operational procedures should be compared?

Philosophers taking a utilitarian approach to ethics frequently apply the notion of utility in moral contexts. These utilitarian philosophers often think of utility as a mental state. That my utility increases if I get a new car means that my mental state is transformed from one state into another, which is more valuable. Let us see if we can make any sense of this traditional utilitarian notion of utility. If we can, we will at least be able to say something interesting about the *meaning* of utility. To start with, it is helpful to divide the utility of an outcome or object into temporal intervals, such that the utility may vary from one interval to the next, but not within an interval. Call such intervals, which may be arbitrarily small, *moments* of utility. It is, of course, an empirical question whether moments of utility exist. It cannot be excluded that in some time periods there are no constant moments of utility. To overcome this problem we assume that if m is an interval which cannot be divided into a sequence of constant intervals, then it is always possible to construct a constant interval m' covering the same time interval, such that $m \sim m'$, by choosing some m' having the right intensity.

A moment of utility is to be thought of as a property of an individual's experience within a certain time interval. The more an agent wants to experience a moment, the higher the utility of the moment. Thus, the agent's well-informed preferences over different moments are likely to be the best way of determining the utility of moments. In this respect the utilitarian concept resembles von Neumann and Morgenstern's approach, since the latter also uses preferences for axiomatizing utility.

Let $M = \{a, b, \ldots\}$ be a set of moments, and let $\succ$ be a relation on M representing strict preference. Indifference is represented by the relation $\sim$. We furthermore suppose that there is a binary operation $\succ$

on M. Intuitively, $a \circ b$ denotes the utility of first experiencing the utility moment a, immediately followed by the utility moment b. The set of utility moments is an *extensive structure* if and only if, for all $a, b; c, d \in M$, the axioms stated below hold true. We recognize the first two axioms from von Neumann and Morgenstern's theory, but note that they now deal with moments rather than lotteries.

Util 1 Either $a \succ b$ or $b \succ a$ or $a \sim b$

Util 2 If $a \succ b$ and $b \succ c$, then $a \succ c$

Util 3 $[a \circ (b \circ c)] \sim [(a \circ b) \circ c]$

Util 4 $a \succ b$ if and only if $(a \circ c) \succ (b \circ c)$ if and only if $(c \circ a) \succ (c \circ b)$

Util 5 If $a \succ b$, then there is a positive integer n such that $(na \circ c) \succ (nb \circ d)$, where na is defined inductively as $1a = a$, $(n + 1)\, a = (a \circ na)$.

The third axiom, Util 3, is mainly technical. It holds that it does not matter if b and c are attached to a or if c is attached to a and b. Util 4 states that in the case that a utility moment a is preferred to b, then $a \succ b$ even in the case that a moment c comes before or after those moments. Of course, this axiom does not imply that the entities that *cause* utility can be attached in this way. For example, if salmon is preferred over beef, it would be a mistake to conclude that salmon followed by ice cream is preferred to beef followed by ice cream. Moreover, Util 4 tells us that if the utility of eating salmon is preferred to the utility of eating beef, then the utility of eating salmon followed by the utility of eating ice cream after eating salmon is preferred to the utility of eating beef followed by the utility of eating ice cream after eating beef.

Util 5 is an Archimedean condition. It implies that even if d is very strongly preferred to c, then this difference can always be outweighed by a sufficiently large number of moments equal to a with respect to b, $a \succ b$, such that $(na \circ c) \succ (ny \circ d)$. This roughly corresponds to the Archimedean property of real numbers: if $b > a > 0$ there exists a finite integer n such that $na > b$, no matter how small a is. Util 5 is problematic if one thinks that there is some critical level of utility, such that a sequence of moments containing a subcritical level moment should never be preferred to a sequence of moments not containing a subcritical level moment. Personally I do not think there are any such critical levels, but to really argue for that point is beyond the scope of this book.

Theorem 5.4 If Util 1–5 hold for $\succ$ on a nonempty set of moments M and if $\circ$ is a binary operation on M, then there exists a real-valued function u on M such that

(1) $a \succ b$ if and only if $u(a) > u(b)$, and
(2) $u(a \circ b) = u(a) + u(b)$, and
(3) For every other function u' that satisfies properties (1) and (2) there exists some $k > 0$ such that $u' = ku$.

Theorem 5.4 follows from a standard theorem in measurement theory. (See e.g. Krantz *et al.* 1971.) I shall spare the reader from the proof. Note that the axioms listed above do not say anything about what is being measured. It is generally agreed that they hold for mass and length, and if the (hedonistic) utilitarians are right they also hold for moments of utility. However, if they hold for moments of utility they merely fix the meaning of the concept. The axioms say almost nothing about how utility could be measured in practice. Are agents really able to state preferences not between, say, salmon and beef, but between the mental states caused by having salmon or beef, respectively? And are they really able to do so even if the comparison is made between hypothetical mental states which are never experienced by anyone, as required by the theory?

Exercises

5.1 Which preference ordering is represented by the following ordinal utility function: $u(a) = 7$, $u(b) = 3$, $u(c) = 34$, $u(d) = -430$, $u(e) = 3.76$?

5.2 The following function u is not an ordinal utility function. Why not? $u(a) = 7$, $u(b) = 3$, $u(c) = 34$, $u(d) = -430$, $u(e) = 3.76$, $u(e) = 12$.

5.3 Your preferences are transitive and asymmetric, and you prefer a to b and b to c. Explain why it has to be the case that you do not prefer c to a.

5.4 Your preferences are asymmetric and complete, and you prefer a to b and b to c. What is your preference between a and c ?

5.5 Show that negative transitivity is logically equivalent with the following claim: $x \succ z$ implies that, for all y in B, $x \succ y$ or $y \succ z$.

5.6 Show that if $\succ$ is asymmetric and negatively transitive, then $\succ$ is transitive.

5.7 You prefer B to A and you are indifferent between receiving A for sure and a lottery that gives you B with a probability of 0.9 and C with a probability of 0.1. You are also indifferent between receiving A for sure and a lottery that gives you B with a probability of 0.6 and D with a probability of 0.4. Finally, you prefer B to A and A to D. All of your preferences satisfy the von Neumann–Morgenstern axioms.
   (a) What do you prefer most, C or D?
   (b) Calculate the (relative) difference in utility between B and C, and between B and D.
   (c) If we stipulate that your utility of B is 1 and your utility of C is 0, what are then your utilities of A and D?

5.8 You are indifferent between receiving A for sure and a lottery that gives you B with a probability of 0.8 and C with a probability of 0.2. You are also indifferent between receiving A for sure and a lottery that gives you B with a probability of 0.5 and D with a probability of 0.5. Finally, you prefer B to A and all of your preferences satisfy the von Neumann–Morgenstern axioms. What is the utility of A, B, C and D (measured on an interval scale of your choice)?

5.9 The continuity axiom employed by von Neumann and Morgenstern holds that if $A \succ B \succ C$ then there exist some probabilities p and q such that $ApC \succ B \succ AqC$. Let $A = \$10{,}000{,}001$ and $B = \$10{,}000{,}000$, and $C =$ 50 years in prison. (a) Do you think it is really true that there are *any* values of p and q such that $ApC \succ B \succ AqC$ truly describes your preferences? (b) Psychological studies suggest that most people cannot distinguish between very small probabilities, i.e. that their preferences over lotteries in which there is a small probability of a very bad outcome are unaffected by exactly how small the probability is. Does this show that there is something wrong with von Neumann and Morgenstern's theory?

5.10 You prefer a fifty-fifty chance of winning either \$100 or \$10 to a lottery in which you win \$200 with a probability of 1/4, \$50 with a probability of 1/4, and \$10 with a probability of 1/2. You also prefer a fifty-fifty chance of winning either \$200 or \$50 to receiving \$100 for sure. Are your preferences consistent with von Neumann and Morgenstern's axioms?

5.11 You somehow know that the probability is 75% that your parents will complain about the mess in your room the next time they see you. What is their utility of complaining?

5.12 You somehow know that the probability is 5% that you will tidy up your room before your parents come and visit you. What is best for you, to tidy up your room or live in a mess?

5.13 The conclusion of Exercise 5.12 may be a bit surprising – can you really find out what is best for you by merely considering what you are likely to do? For example, the probability is 0.99 that a smoker will smoke another cigarette, but it seems false to conclude that it would be better for the smoker to smoke yet another cigarette. What could the advocate of the probabilistic theory say in response to this objection?

Solutions

5.1 $c \succ a \succ e \succ b \succ d$

5.2 Because u assigns two different numbers to one and the same argument: $u(e) = 3.76$ and $u(e) = 12$.

5.3 It follows from transitivity that you prefer a to c, and by applying asymmetry to that preference we find that it is not the case that c is preferred to a.

5.4 We cannot tell what your preference between c and a is, except that it is complete.

5.5 By contraposition, the rightmost part of the statement can be transformed into the following logically equivalent statement: not ($x \succ y$ or $y \succ z$ for all y in B) implies not $x \succ z$. This is equivalent with: for all y in B, not $x \succ y$ and not $y \succ z$ implies not $x \succ z$. This is negative transitivity.

5.6 Suppose that $x \succ y$ and $y \succ z$. From asymmetry and Exercise 5.6 we can conclude that: $x \succ y$ implies that either $x \succ z$ or $z \succ y$. The second possibility, $z \succ y$, is inconsistent with what we initially supposed and can hence be ruled out. Hence, $x \succ z$, which gives us transitivity. Asymmetry implies that not $y \succ x$ and not $z \succ y$. By applying negative transitivity to these preferences we get: not $z \succ x$. From completeness it follows that either $x \succ z$ or $x \sim z$.

5.7 (a) D. (b) The difference between B and C is exactly four times the difference between B and D. (c) $u(A) = 0.9$ and $u(D) = 0.75$.

5.8 $u(B) = 100$, $u(A) = 80$, $u(D) = 60$ and $u(C) = 0$

5.9 (a) No matter how small the probabilities p and q are it seems better to take B for sure. Why risk everything for one extra dollar? (b) The psychological evidence suggests that the von Neumann–Morgenstern theory does not accurately describe how people do in fact behave. Whether it is a valid normative hypothesis depends on what one thinks about the answer to (a).

5.10 No. Your preferences violate the independence axiom.

5.11 Their utility of complaining is three times that of not complaining.

5.12 It is better for you to leave your room as it is; the utility of that option is 19 times that of tidying up the room.

5.13 In this context "better for" means "better as viewed from the decision maker's present viewpoint and given her present beliefs and desires." The point that the smoker may in fact die of lung cancer at some point in the future is perfectly consistent with this notion of "better."