Chapter 2

Decision Problems

In this chapter we will discuss *one-person decision problems.* We start with decision problems under *certainty*, where the decision maker perfectly knows the consequence of each of his possible decisions. We investigate under which circumstances the decision maker's preferences over his decisions allow for an optimal choice, and when his preferences can be represented by a *utility function.*

Afterwards, we turn to decision problems under *uncertainty,* where the consequences of some decisions depend on events that are beyond the decision maker's control, and about which the decision maker is thus uncertain. In this case, the decision maker's preferences over the possible decisions will depend on how likely he deems each of these uncertain events. When taken together, these subjective likelihoods constitute the decision maker's *probabilistic belief.* A *conditional preference relation* then specifies the decision maker's preferences over the decisions for every possible belief that he could have. Throughout this book we will focus on a special class of conditional preference relations: those that have an *expected utility* representation. We explain what we mean by an expected utility representation, and how it can be computed for a given conditional preference relation. At the end of the chapter we characterize, in a convenient way, those choices that are optimal for some belief.

In Chapter 2 of the online appendix to this book, we investigate which conditions must be imposed on a conditional preference relation such that it allows for an expected utility representation. In that same chapter we also discuss some economic applications of the theory we present here.

2.1 Decision Making under Certainty

As announced above, we will start with decision problems under *certainty,* where the decision maker (DM from now on) can fully predict the consequence of each of the choices he can possibly make. The main ingredient of such decision problems will be the *preference relation* that the DM has over the available choices. In particular, we will discuss when such a preference relation allows for *optimal choices,* and when it can be represented by a *utility function.*

2.1.1 Preference Relations

Decision making is an important part of our life. Both at a professional level and in our everyday life we must constantly make choices: Big choices with important consequences, but also many smaller choices with a lower impact. It is fair to say that an important part of our life is shaped by the decisions we make. In this book we focus mostly on decision problems from everyday life which, we hope, most readers will be able to identify with. The leading example for this chapter is "The birthday party", and the story is as follows.

It is Saturday morning, the sun is shining, and later today you will celebrate your birthday. The problem is: Where will you celebrate the party? In your *house*, in the *garden*, or in a *tent* that can be placed in front of your house? From the weather forecast it is clear that it will remain sunny today. Suppose your preferences over the three possible choices are given by

$$\textit{garden} \succ \textit{house}, \quad \textit{tent} \succ \textit{house}, \quad \textit{garden} \succ \textit{tent}. \tag{2.1.1}$$

This should be read as follows: You prefer to have the party in the *garden* rather than having it in the *house*, you prefer to have the party in a *tent* rather than having it in your *house*, and you prefer to have the party in the *garden* rather than having it in a *tent*. Or, equivalently, if you were given the choice between a party in the *garden* and a party in the *house,* you would definitely choose the *garden*, and similarly for the other two pairwise comparisons. Such pairwise comparisions, when taken together, constitute a *preference relation.*

Definition 2.1.1 (Preference relation) *Let C be a finite set of choices for the DM. A* ***preference relation*** $\succsim$ *on C specifies for every pair of choices a, b in C whether (a) the DM prefers a to b, written as $a \succ b$, (b) the DM prefers b to a, written as $b \succ a$, or (c) the DM is indifferent between a and b, written as $a \sim b$.*

Moreover, exactly one of these three cases must hold for every pair of choices a, b. That is, we assume that the DM can evaluate *every* pair of choices *unambiguously* in terms of preference or indifference. For instance, if the DM prefers a to b, than he cannot at the same time prefer b to a, or be indifferent between a and b. The DM would be indifferent between a and b if he does not prefer any of the two options over the other.

We write $a \succsim b$ if the DM either prefers a to b (that is, $a \succ b$), or is indifferent between a and b (that is, $a \sim b$). In this case, we say that the DM *weakly prefers* a to b.

In the example above, you would be the DM and your set of choices would be $C = \{$*house, garden, tent*$\}$. It may be verified that the pairwise comparisons as stated above constitute a preference relation $\succsim$.

2.1.2 Optimal Choices

Let us return to the example "The birthday party", with the preference relation as stated above. Given these preferences, which choice would you make? Clearly, your optimal choice would be to organize the party in your *garden*, as you prefer *garden* to the other two options. In general, a choice is called *optimal* if you weakly prefer that choice to any other available choice.

Definition 2.1.2 (Optimal choice) *A choice a is* ***optimal*** *if $a \succsim b$ for every other choice b.*

Note that an optimal choice need not be (strictly) preferred to every other option – *weakly* preferred is enough. In other words, you may deem an optimal choice *equally good* as some other choice. This will become clear from the following question.

Question 2.1.1 *Suppose that in the example "The birthday party" your preference relation would be as follows: garden $\succ$ house, garden $\sim$ tent, and tent $\succ$ house. What would be your optimal choice(s)?*

A natural question is: Does every preference relation allow for an optimal choice? The answer is "no". To see this, consider again "The birthday party" example, and assume that your preference relation has changed to

$$\textit{garden} \succ \textit{house}, \quad \textit{house} \succ \textit{tent}, \quad \textit{tent} \succ \textit{garden}.$$

Then, there is no optimal choice for you. Indeed, for every choice there is another choice that is preferred to it.

There is something strange about this preference relation, however. If you prefer *garden* to *house* and *house* to *tent* then it seems natural that you also prefer *garden* to *tent.* Your preference relation, on the other hand, indicates that you prefer *tent* to *garden,* which seems contradictory. In formal terms, the preference relation above violates the principle of *transitivity.*

Definition 2.1.3 (Transitivity) *A preference relation $\succsim$ is* ***transitive*** *if for every three choices a, b and c where $a \succsim b$ and $b \succsim c$, it holds that $a \succsim c$.*

The following question collects some important properties that follow from transitivity.

Question 2.1.2 *Consider a transitive preference relation $\succsim$. Show that the following properties hold:*
(a) if $a \succ b$ and $b \succ c$ then $a \succ c$;
(b) if $a \succ b$ and $b \sim c$ then $a \succ c$;
(c) if $a \sim b$ and $b \succ c$ then $a \succ c$;
(d) if $a \sim b$ and $b \sim c$ then $a \sim c$.

It can be shown that under transitivity, an optimal choice will always exist.

Theorem 2.1.1 (Existence of optimal choice) *For every transitive preference relation on a finite set of choices there is at least one optimal choice.*

Here is a proof: Take a finite set C of choices, and a preference relation $\succsim$ on C which is transitive. Suppose, contrary to what we want to show, that there is no optimal choice in C. Then, for every choice $c \in C$ there must be some choice $b(c) \in C$ such that $b(c) \succ c$. Start from an arbitrary choice c^1, and let $c^2 := b(c^1)$, let $c^3 := b(c^2)$, and so on. This way we create an infinite sequence of choices $(c^1, c^2, c^3, ...)$ with the property that $c^{k+1} \succ c^k$ for every k. As there are only finitely many choices in C, there must be a choice that occurs at least twice in this sequence. Let c^* be such a choice, and let k be the first number for which $c^* = c^k$. As c^* occurs at least twice, there must be some $m \geq 1$ with $c^* = c^{k+m}$. We thus see that

$$c^* = c^{k+m} \succ c^{k+m-1} \succ c^{k+m-2} \succ ... \succ c^{k+1} \succ c^k = c^*.$$

As $c^* = c^{k+m} \succ c^{k+m-1} \succ c^{k+m-2}$, it follows from transitivity and property (a) in Question 2.1.2 that $c^* \succ c^{k+m-2}$. As $c^* \succ c^{k+m-2} \succ c^{k+m-3}$ it follows in a similar way that $c^* \succ c^{k+m-3}$. By continuing in this fashion we eventually conclude that $c^* \succ c^{k+1} \succ c^k = c^*$. Hence, $c^* \succ c^{k+1}$ and $c^{k+1} \succ c^*$, which is a contradiction. Thus, there must be an optimal choice. This completes the proof. ■

In the light of Theorem 2.1.1, one may wonder whether every preference relation that admits an optimal choice must be transitive. The following question reveals that the answer is "no": There are intransitive preference relations that still allow for optimal choices.

Question 2.1.3 *Consider the set $C = \{a, b, c, d\}$ of choices, and the preference relation $\succsim$ given by*

$$b \sim a, \quad a \sim c, \quad d \succ a, \quad b \succ c, \quad b \succ d, \quad c \succ d.$$

(a) Show that there is an optimal choice.
(b) Explain why $\succsim$ is not transitive.

2.1.3 Utility Representations

Consider the example "The birthday party" with the preference relation given by (2.1.1). That is, you prefer *garden* to *tent* and *tent* to *house,* and hence, by transitivity, you also prefer *garden* to *house.* Such preferences may be represented by the following *utilities*:

$$u(garden) = 3, \quad u(tent) = 2, \quad u(house) = 1. \tag{2.1.2}$$

This should be read as follows: Having the party in the *garden* gives you a utility of 3, having the party in a *tent* gives you a utility of 2, and having the party in your *house* gives you a utility of 1. These utilities indicate how *good* or *desirable* you deem each of the three options, relative to the other options.

The preferences in (2.1.1) can, in fact, be derived from the utility function in (2.1.2): If we compare *garden* and *tent* then we see that *garden* has the higher utility, and hence you prefer *garden* to *tent.* By analyzing the other two pairs of choices, we similarly see that in (2.1.1) you always prefer the choice that yields the higher utility in (2.1.2). In that sense, the utility function in (2.1.2) *represents* the preference relation in (2.1.1).

Definition 2.1.4 (Utility representation) *Consider a set of choices C. A* ***utility function*** *u assigns to every choice $c \in C$ some number $u(c)$. The utility function u* ***represents*** *a preference relation $\succsim$ if for every pair of choices a, b we have that*

$$a \succsim b \text{ if and only if } u(a) \geq u(b).$$

There are many other utility representations for the preference relation in (2.1.1). For instance, the utility functions v and w given by

$$v(garden) = 8, \quad v(tent) = 7, \quad v(house) = 0 \text{ and}$$
$$w(garden) = 5, \quad w(tent) = 0, \quad w(house) = -1$$

also represent this preference relation. Hence, the utilities, when viewed in isolation, have no intrinsic meaning. When we say that the utility of having the party in the tent is 7, then this does not tell us how desirable you find this option. Only the utilities in *comparision* with the other utilities have a meaning: When comparing two choices, you always prefer the choice that yields the higher utility among these two.

Note that in the utility function v above, the utility difference between *garden* and *tent* is 1, whereas the utility difference between *tent* and *house* is 7. This might suggest that the *intensity* by which you prefer *tent* to *house* is higher than the intensity by which you prefer *garden* to *tent.* This is not necessarily the case, however. Look at the utility function w above. There, the utility difference between *garden* and *tent* is 5, whereas the utility difference between *tent* and *house* is 1, but it represents exactly the same preference relation. Hence, also the relative utility differences do not have an intrinsic meaning here. This will be different in the following sections when we investigate decision making under *uncertainty.* As we will see, the relative utility differences in those settings *will* typically reflect the DM's preference intensity.

Question 2.1.4 *Consider the set of choices* $C = \{a, b, c, d\}$, *and the utility function* u *where* $u(a) = 2$, $u(b) = 5$, $u(c) = 7$, $u(d) = 5$. *Describe the preference relation that is represented by the utility function* u.

A practical advantage of a utility representation is that it is much more compact than describing the preference relation in the original way, by means of pairwise rankings. The difference is especially large when there are many choices present. Suppose there are 10 choices. Then, a preference relation must specify 45 pairwise rankings between choices, whereas a utility function only needs to specify 10 numbers. Which of the two representations do you prefer? Moreover, once we have a utility representation it becomes very easy to identify the optimal choices: These are precisely the choices that yield the highest utility.

In this light, it is important to know when a preference relation has a utility representation. The answer is surprisingly simple: Precisely when the preference relation is transitive.

Theorem 2.1.2 (When a utility representation exists) *Consider a finite set of choices* C. *A preference relation* $\succsim$ *on* C *has a utility representation, if and only if,* $\succsim$ *is transitive.*

Once we rely on Theorem 2.1.1, the proof is not so difficult. We need to prove two directions here. Let us first assume that $\succsim$ has a utility representation, that is, $\succsim$ is represented by a utility function u.

Question 2.1.5 *Show that* $\succsim$ *is transitive.*

To show the other direction, assume that $\succsim$ is transitive. We will show that $\succsim$ can be represented by a utility function u. By Theorem 2.1.1 we know that there is at least one optimal choice. Let C^1 be the set of optimal choices.

Question 2.1.6 *(a) Show that* $a \sim b$ *for all* $a, b \in C^1$.
(b) Show that $a \succ b$ *for all* $a \in C^1$ *and for all* $b \notin C^1$.

Now, focus on the set of choices $C \backslash C^1$, that is, all choices that are not in C^1. Consider the restriction of the preference relation $\succsim$ to choices in $C \backslash C^1$. Then, this restricted preference relation is again transitive. By Theorem 2.1.1 we know that there is at least one optimal choice in $C \backslash C^1$ for the restricted preference relation. Let C^2 be the set of optimal choices within $C \backslash C^1$ for the restricted preference relation. In a similar way as in Question 2.1.6 it can then be shown that $a \sim b$ for all $a, b \in C^2$, and $a \succ b$ for all choices $a \in C^2$ and all choices b that are not in $C^1 \cup C^2$. If we continue in this fashion, we finally arrive at sets of choices $C^1, C^2, ..., C^K$ where

$$a \sim b \text{ for all } a, b \text{ in the same set } C^k, \text{ and} \tag{2.1.3}$$

$$a \succ b \text{ whenever } a \in C^k \text{ and } b \in C^{k+1} \cup ... \cup C^K. \tag{2.1.4}$$

We are now ready to define the utility function u. Choose numbers $u^1 > u^2 > ... > u^K$, and set $u(a) := u^k$ whenever $a \in C^k$. Then, in the light of (2.1.3) and (2.1.4), the utility function u represents the preference relation $\succsim$. This completes the proof. ■

Question 2.1.7 *Consider the set of choices* $C = \{a, b, c, d\}$, *and the preference relation* $\succsim$ *given by* $b \succ a$, $c \succ a$, $a \succ d$, $b \sim c$, $b \succ d$, *and* $c \succ d$. *Show that* $\succsim$ *is transitive, and find a utility representation for* $\succsim$.

In the remainder of this book we will solely focus on preference relations that are transitive.

	rainy	stormy	calm
house	guests are happy to be inside	guests are happy to be inside	guests regret being inside
garden	guests become wet and angry	guests have difficulty standing upright	guests enjoy the lovely weather at the fullest
tent	guests remain dry, but are annoyed by the rain ticking on the tent	complete disaster, as the tent will blow away	guests somewhat enjoy the sunny weather, but would rather be outside

Table 2.2.1 Consequences of your decisions, contingent on the weather

2.2 Decision Making under Uncertainty

In the previous section we assumed that the DM was certain about the consequences of each of his decisions. For that reason, we refer to such situations as decision making under *certainty.* From this section onwards we focus on situations where the DM is *uncertain* about the consequences of some of his choices, because the consequence may depend on events that are beyond his control, like the state of the weather, the outcome of some boxing match, or the decisions made by other people. We will see that such uncertainty about the consequences may be modelled by the introduction of *states*, and the assumption that the DM holds a probabilistic *belief* about the states. In this scenario, the DM's preference relation over the choices will be contingent on the belief he holds about the states, and such situations are called *decision problems under uncertainty.* The object that summarizes the DM's contingent preferences for every possible belief about the states is called a *conditional preference relation.* In the remainder of this chapter, we will be interested in those conditional preference relations that can be represented by a *utility matrix,* assigning a utility to every possible choice *and state.* In this case we say that the conditional preference relation has an *expected utility* representation.

2.2.1 States

Consider the example "The birthday party" from the previous section. So far we assumed that you could perfectly predict the consequences of each of the three possible choices: *garden, tent* and *house.* This amounts to saying that you were certain about the weather for that day. If, for instance, you are certain that it will remain sunny and dry during the whole day, then the consequence of having the party in the *garden* will be that the guests will be enjoying the lovely wheather outside, whereas having the party in the *house* would disappoint the guests, because they will be wondering why they cannot have their drink outside with such splendid weather.

Let us now assume that the weather forecasts for today are highly uncertain. There is a chance that the weather will be calm and sunny, but it could also start to rain, and there is even a chance that it will be stormy today. Clearly, the consequences of each of the possible decisions will crucially depend on which of these three weather scenarios will be realized. And, as a consequence, your preferences over the choices will depend on the weather scenario as well. Table 2.2.1 gives a possible description of the consequences of each of the three decisions, contingent on the state of the weather.

As a result, your preferences over the three choices, contingent on the state of the weather, would be given by

$$\begin{aligned} &\textit{house} \succ_{[r]} \textit{tent} \succ_{[r]} \textit{garden}, \\ &\textit{house} \succ_{[s]} \textit{garden} \succ_{[s]} \textit{tent}, \\ &\textit{garden} \succ_{[c]} \textit{tent} \succ_{[c]} \textit{house}. \end{aligned} \tag{2.2.1}$$

Here, $\succsim_{[r]}, \succsim_{[s]}$ and $\succsim_{[c]}$ denote your preference relations if the weather is *rainy, stormy,* or *calm,* respectively. Moreover, we assume that these preference relations are all transitive. Thus, when we write *house* $\succ_{[r]}$*tent* $\succ_{[r]}$*garden,* we mean that *house* $\succ_{[r]}$*tent,* that *tent* $\succ_{[r]}$*garden,* and therefore, by transitivity, *house* $\succ_{[r]}$*garden.*

The three possible weather conditions, *rainy, stormy* and *calm,* are called *states.* In general, a state represents an event that influences the consequence of one or more decisions, and which is beyond the DM's control. As such, the state will also determine how the DM's preferences over his choices will look like. This is clearly shown by (2.2.1), where your preferences over the three possible party locations depends on the state (of the weather). A general requirement is that the states represent events that are *mutually exclusive* and *exhaustive,* that is, at all times exactly one of these events must be the true event.

2.2.2 Beliefs

The question is now: Which of the three party locations would you choose in the birthday example? The only correct answer is: "It depends". Indeed, it depends on what you expect the weather to be in the afternoon and evening. If you expect the weather to be *rainy,* or expect the weather to be *stormy,* the optimal location would be your *house.* If instead you expect the weather to be *calm*, then it would be optimal to have the party in the *garden.*

Does this mean that it will never be optimal for you to have the party in the *tent*? Well, not necessarily. So far, we have only considered beliefs about the weather where you assign *probability one* to one of the three states. When we said above that "you expect that the weather will be *rainy*", what we are really saying is that you assign probability 1 to the event that it will rain. Or, in other words, you are 100% sure that it will rain.

But typically our beliefs are more modest than this. If you are truly uncertain about the weather, then it is reasonable to assign a substantial probability to each of the three states. For instance, you could assign probability 0.3 to the weather being *rainy,* probability 0.1 to the weather being *stormy,* and probability 0.6 to the weather being *calm.* This reflects a state of mind in which you deem each of the three weather scenarios possible, but where you deem the state *calm* twice as likely as the state *rainy,* and deem the state *rainy* three times as likely as the state *stormy.* This is called a *probabilistic belief,* or simply a *belief.*

Definition 2.2.1 (Probability distribution and belief) *Consider a finite set X. A* ***probability distribution*** *on X is a function p that assigns to every $x \in X$ a number $p(x)$, such that $p(x) \geq 0$ for all $x \in X$, and $\sum_{x \in X} p(x) = 1$. The set of all probability distributions on X is called $\Delta(X)$.*

Consider a finite set of states S. A ***belief*** *on S is a probability distribution $p \in \Delta(S)$ on S.*

If we consider a belief p on S then, for a given state s, the number $p(s)$ reflects how *likely* the DM deems this particular state, compared to the other states.

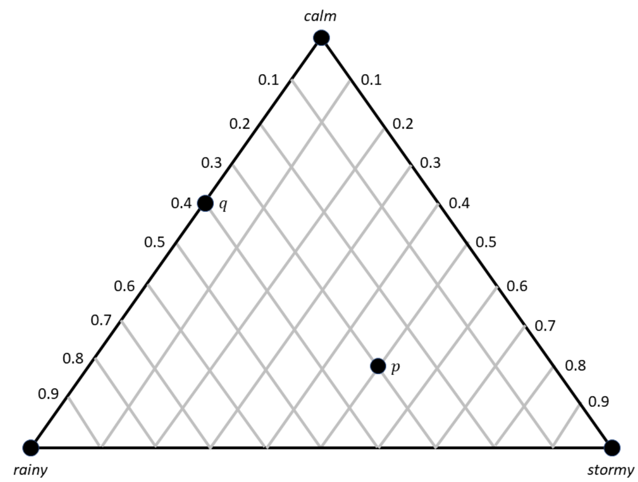

Figure 2.2.1 Visualization of beliefs with three states

2.2.3 Conditional Preference Relations

If we allow for such probabilistic beliefs over states, then it is natural to assume that your preferences over the three party locations will depend not only on the three possible weather scenarios, but on your *belief* about these weather scenarios. For instance, if you assign probability 0.3 to the weather being *rainy,* probability 0.1 to the weather being *stormy,* and probability 0.6 to the weather being *calm,* as above, then your preferences over the three party locations may be different than when you would assign probability 0.5 to *rainy,* probability 0.2 to *stormy,* and probability 0.3 to *calm.* And these two preference relations may, in turn, be different from a scenario where you assign probability 0 to *rainy,* probability 0.6 to *stormy,* and probability 0.4 to *calm.* The object that summarizes the DM's preferences for every possible belief is called a *conditional preference relation.*

Definition 2.2.2 (Conditional preference relation) *Consider a set of choices C and a set of states S. A* ***conditional preference relation*** *$\succsim$ assigns to every belief $p \in \Delta(S)$ a preference relation $\succsim_p$ over the choices in C.*

If we put together the set of choices, the set of states, and the conditional preference relation, we obtain a *decision problem under uncertainty.*

Definition 2.2.3 (Decision problem under uncertainty) *A* ***decision problem under uncertainty*** *is a triple $(C, S, \succsim)$, where C is the set of choices, S is the set of states, and $\succsim$ is a conditional preference relation.*

In the remainder of this book we will refer to such objects simply as *decision problems,* as we will always focus on situations with uncertainty. The main ingredient of a decision problem is thus the conditional preference relation, assigning to every possible belief a preference relation over the possible choices. If there are two, three or four states, then such conditional preference relations can be represented *graphically,* as we will see in this chapter.

Consider again the example "The birthday party", where the set of states is $S = \{$*rainy, stormy, calm*$\}$. Then, every belief $p \in \Delta(S)$ can be identified with a point in a triangle, as depicted in Figure 2.2.1. The three cornerpoints of the triangle correspond to the three beliefs where you assign

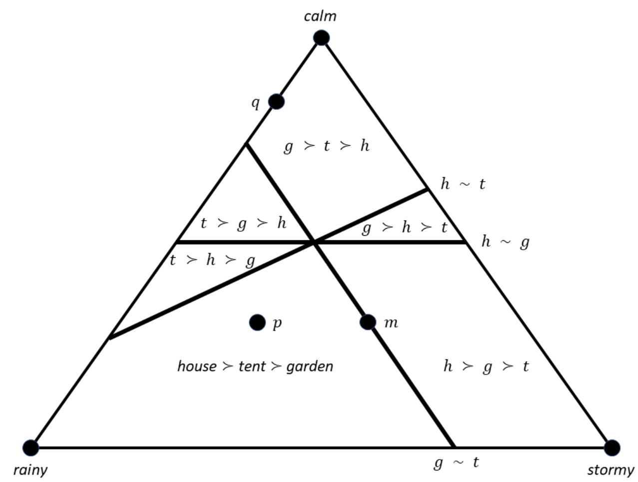

Figure 2.2.2 Conditional preference relation for "The birthday party"

probability 1 to *rainy,* probability 1 to *stormy,* and probability 1 to *calm,* respectively. On the edge between *rainy* and *calm* we have denoted the probability that the belief assigns to *rainy.* Consider, for instance, the probability 0.7 on this edge. Then, the grey line that starts at 0.7 represents all beliefs that assign probability 0.7 to *rainy.* Similarly for the other grey lines that start at a probability on this edge. In the same fashion, the probabilities at the edge between *stormy* and *calm* denote the probability that the belief assigns to *stormy.* The grey line that starts at probability 0.2 on this edge represents all beliefs that assign probability 0.2 to *stormy,* and similarly for the other grey lines that start at a probability on this edge.

Now, consider the belief p in Figure 2.2.1. By looking at the two grey lines that pass through p, it may be verified that p assigns probability 0.3 to *rainy* and probability 0.5 to *stormy.* As the sum of the probabilities must be 1, the belief p must assign probability 0.2 to *calm.* We will write $p = (0.3, 0.5, 0.2)$, where these three numbers represent the probabilities assigned to the first state (*rainy*), the second state (*stormy*) and the third state (*calm*), respectively. Alternatively, the belief p can also be written as $p = (0.3)\cdot$ *rainy* $+ (0.5)\cdot$ *stormy* $+ (0.2)\cdot$ *calm.* We use both types of notation, depending on what is most convenient.

Question 2.2.1 *(a) Describe the belief q in Figure 2.2.1.*
(b) Locate the belief $r = (0.6, 0.3, 0.1)$ in the triangle of Figure 2.2.1.

Now that we know how your beliefs can be identified with points in a triangle, we can can accordingly visualize a conditional preference relation. Consider, for instance, the conditional preference relation depicted in Figure 2.2.2. As you can see, the set of all possible beliefs is divided into six different main areas, and for each of the main areas we have specified the preferences over choices you would have there. For instance, at the belief p your preferences would be

$$\textit{house} \succ_p \textit{tent} \succ_p \textit{garden},$$

whereas at the belief q your preferences would be given by

$$\textit{garden} \succ_q \textit{tent} \succ_q \textit{house}.$$

Then, there are also the three lines where you would be indifferent between two of the choices. For instance, the line indicated by $g \sim t$ represents the beliefs where you would be indifferent between *garden* and *tent.* At the belief m, depicted in the figure on this line, your preferences would thus be given by

$$\textit{house} \succ_m \textit{garden} \sim_m \textit{tent.}$$

Note that this conditional preference relation is consistent with the preferences contingent on the three states in (2.2.1). Indeed, at the three cornerpoints of the triangle, the preferences are given by

$$\begin{aligned} &\textit{house} \succ_{[r]} \textit{tent} \succ_{[r]} \textit{garden,} \\ &\textit{house} \succ_{[s]} \textit{garden} \succ_{[s]} \textit{tent,} \\ &\textit{garden} \succ_{[c]} \textit{tent} \succ_{[c]} \textit{house,} \end{aligned}$$

where $[r], [s]$ and $[c]$ are the three beliefs that assign probability 1 to *rainy, stormy* and *calm,* respectively. At these beliefs, where you are 100% certain about the weather condition, it is never optimal to have the party in a *tent.*

At the time, it can be seen from Figure 2.2.2 that choosing the *tent* would be optimal for beliefs that assign positive probability to both *rainy* and *calm.* That is, if you are truly uncertain about whether it will rain or not, it may be optimal to go for the *tent.* This also makes intuitive sense: Although *tent* is not optimal at the states *rainy* and *calm,* it provides a good second-best option for both states. Moreover, organizing the party in the *house* or in the *garden* is more risky, as both options could lead to really bad consequences in one of these two states. Hence, there is something to say for the "safe" option *tent* if you are highly uncertain about the chances of a rain shower. As the choice *tent* is optimal for at least one belief, we call *tent* a *rational* choice.

Definition 2.2.4 (Rational choice) *A choice $c \in C$ is* ***rational*** *if there is a belief $p \in \Delta(S)$ such that c is optimal for the preference relation $\succsim_p$.*

Given the preferences at the three states as specified in (2.2.1) it is not guaranteed, however, that *tent* is a rational choice. To see this, consider the alternative conditional preference relation in Figure 2.2.3. This conditional preference relation still induces the same preference relations at the three states, as given by (2.2.1). However, compared to the conditional preference relation from Figure 2.2.2, your preference intensity for the choice *tent* has gone down. This can be seen from the fact that the set of beliefs for which you prefer *tent* to *house,* and the set of beliefs for which you prefer *tent* to *garden,* have both shrunk relative to Figure 2.2.2. It could be that you suddenly remember the terrible camping experience you had last summer, decreasing the appeal of a *tent.*

2.2.4 Dominance and Preference Reversals

Consider a conditional preference relation $\succsim$, and two choices a and b. The following definition reveals what it means that choices a and b are *equivalent*, that *a strictly dominates b*, that *a weakly dominates b*, and that there are *preference reversals* between a and b.

Definition 2.2.5 (Equivalence, dominance, and preference reversals) *Consider a conditional preference relation $\succsim$ and two choices a and b.*

(a) Choices a and b are ***equivalent*** *if $a \sim_p b$ for all beliefs p.*

(b) Choice a ***strictly dominates*** *choice b if $a \succ_p b$ for all beliefs p.*

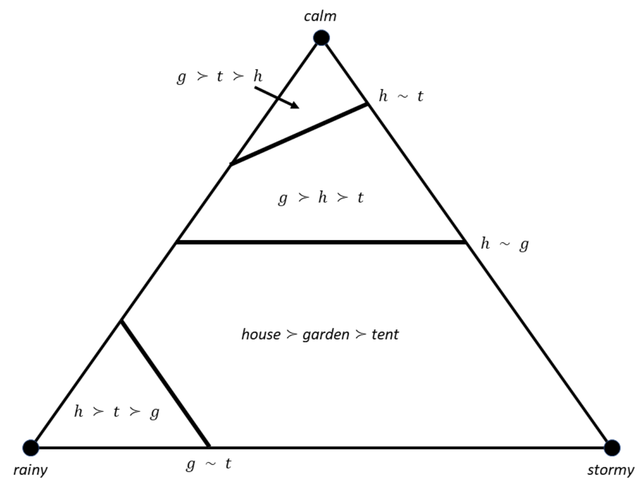

Figure 2.2.3 Alternative conditional preference relation for "The birthday party"

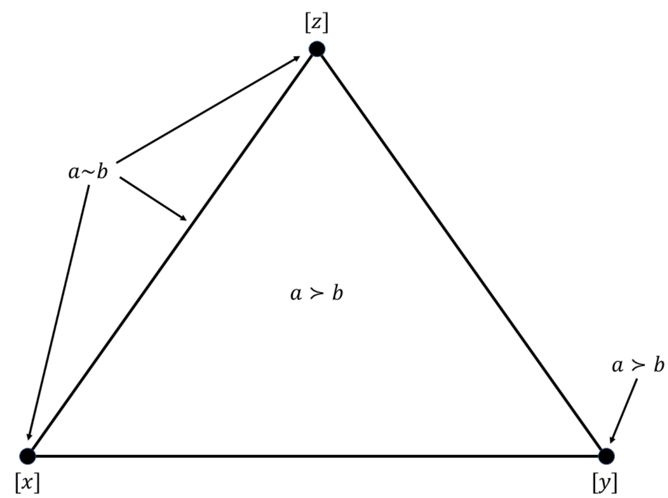

Figure 2.2.4 Illustration of weak dominance

*(c) Choice a **weakly dominates** choice b if $a \succsim_p b$ for all beliefs p, and $a \succ_q b$ for some belief q.*

*(d) There are **preference reversals** between a and b if there is a belief p with $a \succ_p b$ and some other belief q with $b \succ_q a$.*

Note that for every two choices a and b, we must either have that (i) a and b are equivalent, or (ii) a weakly dominates b, or (iii) b weakly dominates a, or (iv) there are preference reversals between a and b.

As an illustration, consider the conditional preference relation in Figure 2.2.4. Here, $[x]$ denotes the belief that assigns probability 1 to state x. Similarly for $[y]$ and $[z]$. Since the DM is indifferent between choices a and b for all beliefs on the line segment between $[x]$ and $[z]$, but strictly prefers a to b for all other beliefs, we conclude that choice a weakly dominates choice b, but a does not strictly dominate b.

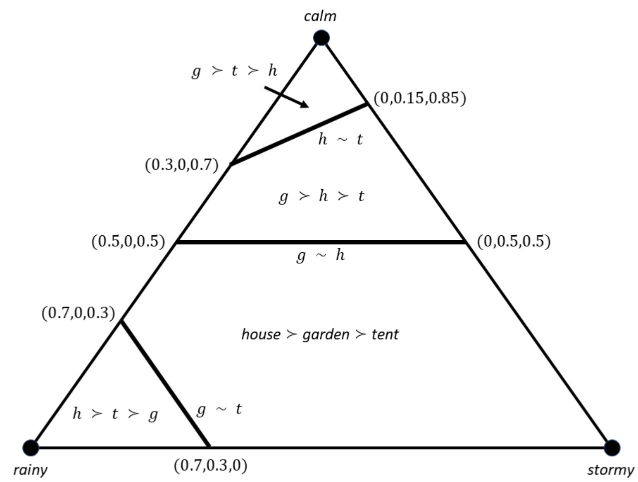

Figure 2.3.1 Conditional preference relation for "The birthday party"

	rainy	*stormy*	*calm*
house	10	17	0
garden	0	7	10
tent	3	0	3

Table 2.3.1 Expected utility representation for the conditional preference relation in Figure 2.3.1

2.3 Expected Utility Representation

In this section we will see what it means that a conditional preference relation is represented by a matrix of utilities. We refer to such a matrix as an *expected utility representation.* Afterwards, we will explore how *unique* the expected utility representation is. Finally, we will interpret the differences in expected utility from two choices a and b as the *intensity* by which the DM prefers a to b.

2.3.1 Expected Utility

Consider the conditional preference relation in Figure 2.3.1. This is the same conditional preference relation as in Figure 2.2.3, but now supplemented with some precise numbers. As we will see, the conditional preference relation is represented by the *utility matrix* in Table 2.3.1. This utility matrix thus specifies a utility for every choice *and every state.* For instance, if you have the party in the *garden* and the weather is *stormy,* then the resulting utility would be 7.

Consider some belief p, say $p = (0.3)\cdot$ *rainy* + $(0.5)\cdot$ *stormy* + $(0.2)\cdot$ *calm.* Under this belief, your *expected utility* from choosing *house* would be

$$u(\textit{house},\ p) = (0.3) \cdot 10 + (0.5) \cdot 17 + (0.2) \cdot 0 = 11.5.$$

The logic behind this expression is that you multiply each of the possible utilities that the choice *house* can induce by the probability you expect this utility to occur, and you sum these terms up. For instance, you expect the weather to be *rainy* with probability 0.3, in which case your utility would be 10. Similarly for the other two terms above.

In the same way we can derive the expected utility of choosing *garden* and *tent* under this belief p, yielding

$$u(\textit{garden}, p) = (0.3)\cdot 0 + (0.5)\cdot 7 + (0.2)\cdot 10 = 5.5, \text{ and}$$
$$u(\textit{tent}, p) = (0.3)\cdot 3 + (0.5)\cdot 0 + (0.2)\cdot 3 = 1.5.$$

In particular, we see that the expected utility of *house* is higher than that of *garden,* which in turn is higher than that of *tent.* This matches precisely your preferences at the belief p which, based on Figure 2.3.1, are *house* $\succ_p$ *garden* $\succ_p$ *tent.* That is, at the belief p you always prefer, among any pair of choices, the choice that yields the higher expected utility.

Question 2.3.1 *Compute the expected utility of the three choices under the belief* $q = (0.3, 0.1, 0.6)$*, and verify that this matches your preferences at* q*.*

Now, consider the belief $r = (0.3)\cdot$ *rainy* $+ (0.7)\cdot$ *calm,* which assigns probability 0 to *stormy.* Then, the expected utilities of the three choices are

$$\begin{aligned} u(\textit{house}, r) &= (0.3)\cdot 10 + 0\cdot 17 + (0.7)\cdot 0 = 3, \\ u(\textit{garden}, r) &= (0.3)\cdot 0 + 0\cdot 7 + (0.7)\cdot 10 = 7, \\ u(\textit{tent}, r) &= (0.3)\cdot 3 + 0\cdot 0 + (0.7)\cdot 3 = 3. \end{aligned}$$

Hence, the expected utility of *garden* is higher than that of the other two choices, whereas *house* and *tent* yield the same expected utility. This matches exactly your preferences at the belief r which, on the basis of Figure 2.3.1, are given by *garden* $\succ_r$ *house* $\sim_r$ *tent.*

It may be verified that, at every belief p, the expected utilities of the three choices always match precisely the preferences you have at p. As an additional exercise, it may be useful to try a few beliefs, to compute the expected utilities of the three choices at these beliefs, and to verify that the induced rankings coincide precisely with your preferences at these beliefs. In this case, we say that the utility matrix *represents* the conditional preference relation.

Definition 2.3.1 (Expected utility representation) *Consider a decision problem* $(C, S, \succsim)$.

*(a) A **utility function** u assigns to every choice $c \in C$ and state $s \in S$ a number $u(c, s)$. This number is called a utility.*

*(b) At a given belief $p \in \Delta(S)$, and for a given choice c, the **expected utility** of choice c under the belief p is*

$$u(c, p) := \sum_{s\in S} p(s)\cdot u(c, s).$$

*(c) The utility function u **represents** the conditional preference relation $\succsim$ if for every belief $p \in \Delta(S)$, and every two choices a, b,*

$$a \succsim_p b \text{ if and only if } u(a, p) \geq u(b, p).$$

*In this case, u is called an **expected utility representation** of the conditional preference relation* $\succsim$.

If the conditional preference relation is represented by a utility function u, we often write the decision problem as (C, S, u), replacing the conditional preference relation by a utility function that represents it.

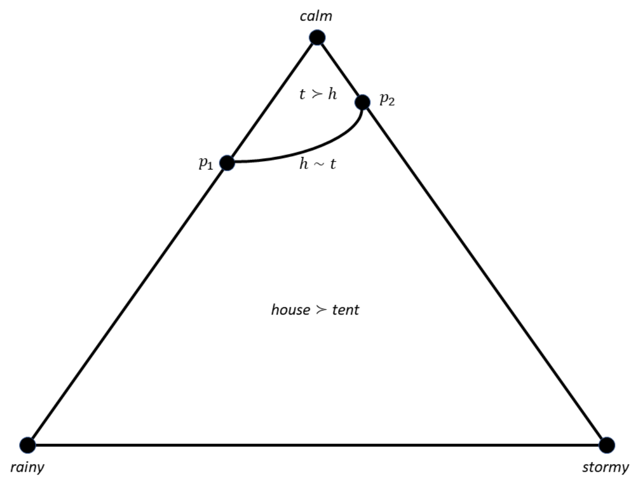

Figure 2.3.2 When there is no expected utility representation

Above we have seen that the conditional preference relation in Figure 2.3.1 has an expected utility representation. But is this is true for *every* conditional preference relation? The answer is "no". As an illustration, consider the conditional preference relation $\succsim$ between the choices *house* and *tent* in Figure 2.3.2. It may be verified that $\succsim$ has no expected utility representation. To see this, suppose that there would be an expected utility representation u for $\succsim$. Then, at the beliefs p_1 and p_2 the expected utility of *tent* should be equal to the expected utility of *house.* This would imply that the expected utilities of *tent* and *house* are equal for every belief on the line between p_1 and p_2. Hence, you would be indifferent between *tent* and *house* for every belief on the line between p_1 and p_2. This, however, contradicts the data from Figure 2.3.2. We thus conclude that there is no expected utility representation for $\succsim$.

This raises the question: What conditions must we impose on a conditional preference relation such that it will have an expected utility representation? This question will be addressed in Chapter 2 of the online appendix to this book. The interested reader is referred to the online appendix for more details. Our general approach there is that we consider three different scenarios: the case of *two choices* only, the case where there are *preference reversals* between every two choices, and the *general* case. For every scenario we present conditions – also called *axioms* – on conditional preference relations that are both necessary and sufficient for allowing an expected utility representation. From Chapter 3 onwards, we will always assume that the conditional preference relations have an expected utility representation.

2.3.2 How Unique is Expected Utility Representation?

Clearly, a utility function u always induces a *unique* conditional preference relation: The conditional preference relation $\succsim$ where for every belief p and every two choices a and b, the DM always prefers the choice that yields the higher expected utility. But does every such conditional preference relation have a *unique* expected utility representation?

The answer is "no". Consider the expected utility representation u in Table 2.3.1, and add 1 utility unit for all choices at the state *rainy,* add 5 utility units for all choices at the state *stormy,* and substract 3 utility units for all choices at the state *calm.* This results in the utility function v as depicted in Table 2.3.2.

	rainy	*stormy*	*calm*
house	11	22	−3
garden	1	12	7
tent	4	5	0

Table 2.3.2 Alternative expected utility representation for the conditional preference relation in Figure 2.3.1

	rainy	*stormy*	*calm*
house	40	68	0
garden	0	28	40
tent	12	0	12

Table 2.3.3 Alternative expected utility representation for the conditional preference relation in Figure 2.3.1

Note that at every state, the utility differences between the three choices are exactly the same as in Table 2.3.1. Consequently, the *expected* utility differences between the three choices at every *belief* will also be the same as in Table 2.3.1. But then, also the new utility function v from Table 2.3.2 will represent the conditional preference relation $\succsim$ from Figure 2.3.1. Indeed, for every belief the DM prefers a choice a to a choice b precisely when the expected utility difference between a and b given the utility function u (and therefore also given the utility function v) is larger than 0.

We can also multiply all utilities in Table 2.3.1 by the same positive number, say 4. This would result in the utility function w as depicted in Table 2.3.3. With the new utility function w, the utility differences at every state are exactly 4 times as large as in Table 2.3.1, and hence the same applies to the expected utility differences at every belief. Therefore, the new utility function w will also represent the conditional preference relation $\succsim$ from Figure 2.3.1. The argument is similar as above: For every belief the DM prefers a choice a to a choice b precisely when the expected utility difference between a and b given the utility function u (and therefore also given the utility function w) is larger than 0.

We thus see that, if we take an expected utility representation u, and multiply all utilities by the same positive number $\alpha > 0$, or add for every state s a constant number v_s (possibly negative) to the utilities $u(c, s)$ of the different choices, then the new utility function will still be an expected utility representation. In particular, the expected utility representation is far from being unique.

But we can say a little more: By the argument above there are at least $|S| + 1$ degrees of freedom for choosing the expected utility representation, provided the DM is not always indifferent between all of his choices. Here, $|S|$ denotes the number of states. One degree of freedom arises because we can choose any $\alpha > 0$, whereas the choices of the numbers v_s above, for every state $s \in S$, results in $|S|$ additional degrees of freedom.

For the conditional preference relation in Figure 2.3.1 we would thus have at least 4 degrees of freedom for the expected utility differences, as there are three states. In fact, one can prove that there are *exactly* 4 degrees of freedom here, no more and no less. This will follow from Theorem 2.5.1 in this chapter.

Question 2.3.2 *Consider the conditional preference relation from Figure 2.3.1. Find the unique expected utility representation u where u(house, rainy) $= 0$, u(house, stormy) $= 0$, u(house, calm) $= 0$ and u(garden, calm) $= 1$.*

2.3.3 Expected Utility Differences as Preference Intensities

Above we have argued that for the conditional preference relation in Figure 2.3.1 there are precisely 4 degrees of freedom when choosing an expected utility representation: One because we can multiply all utilities in Table 2.3.1 by the same positive number $\alpha > 0$, and three more because for every state s in Table 2.3.1 we can add the same number v_s to all utilities in the column that belongs to s. But then, the utility differences between the three choices at the three states will be *unique* up to the positive multiplicative constant α, and the same is true for the *expected* utility differences at the various beliefs.

In other words, the *relative* expected utility differences are *unique* at every belief. These relative expected utility differences at the various beliefs have a very natural interpretation, as they may be viewed as the relative *intensities* by which the DM prefers one choice to another at these beliefs. As an illustration, consider the expected utility representation in Table 2.3.1, and take the belief $p = (0.3)\cdot$ *rainy* + $(0.5)\cdot$ *stormy* + $(0.2)\cdot$ *calm.* As we have seen above, the associated expected utilities are

$$u(\textit{house, } p) = 11.5,\ u(\textit{garden, } p) = 5.5 \text{ and } u(\textit{tent, } p) = 1.5.$$

Thus, the expected utility difference between *house* and *garden* is 6, whereas the expected utility difference between *garden* and *tent* is 4. This indicates that the intensity by which the DM prefers *house* to *garden* is 6/4 times as high as the intensity by which he prefers *garden* to *tent.* In other words, at the belief p the DM's preference of *house* over *garden* is 1.5 times stronger than his preference of *garden* over *tent.*

Or consider the belief q that assigns probability 1 to the state *rainy,* where the (expected) utilities are given by

$$u(\textit{house, } q) = 10,\ u(\textit{garden, } q) = 0 \text{ and } u(\textit{tent, } q) = 3.$$

This reflects a state of mind in which the intensity by which the DM prefers *house* to *tent* is 7/3 times as high as the intensity by which he prefers *tent* to *garden.*

We can even compare preference intensities across different beliefs. By comparing the expected utility differences at the two beliefs p and q, we can say that the intensity by which the DM prefers *house* to *garden* at the belief q is 10/6 times higher than at the belief p. Similarly, the intensity by which the DM prefers *house* to *garden* at the belief q is 10/4 times higher than the intensity by which the DM prefers *garden* to *tent* at the belief p.

Not only does an expected utility representation allow us to derive the DM's *preference intensities* at the various beliefs, which is important conceptually, it also has a big practical advantage: Instead of having to specify the DM's preferences over choices for each of the (infinitely many) beliefs, the only thing we need to do is to write down a few utilities, one for every choice-state pair. Especially when there are more than four states, in which case drawing the conditional preference relation becomes either very difficult or impossible, such an expected utility representation would still enable us to summarize the conditional preference relation in a compact way, by means of a matrix with finitely many utilities.

2.4 Utility Design Procedure

Suppose we consider a conditional preference relation that has an expected utility representation. How can we *compute* such an expected utility representation then? This is the question we will address in this section. Indeed, we will develop a *utility design procedure* that allows us to compute an expected

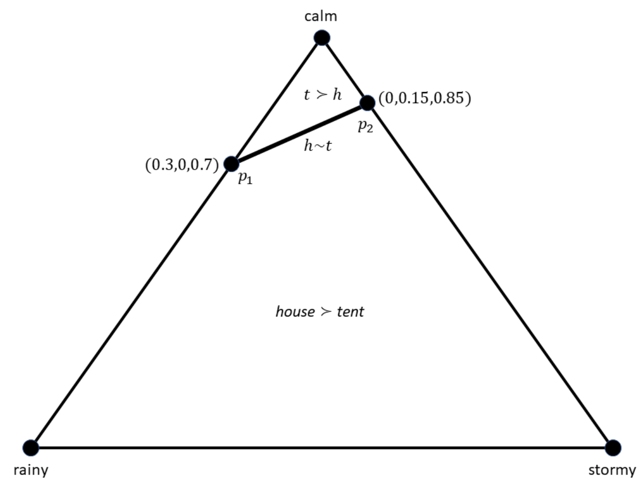

Figure 2.4.1 Utility design procedure for two choices

utility representation in an efficient manner. Since we would like to build up the complexity step by step, we consider three different scenarios with increasing degrees of difficulty: (1) the case of two choices with preference reversals, (2) the case of more than two choices with preference reversals, and (3) the general case.

2.4.1 Two Choices with Preference Reversals

Let us return to the example "The birthday party". Suppose the DM holds the conditional preference relation in Figure 2.3.1. If we focus only on the choices *house* and *tent* we obtain the conditional preference relation in Figure 2.4.1. It turns out that this conditional preference relation has an expected utility representation. But how can we find such an expected utility representation in a systematic way?

Since only the utility *differences* matter, we can choose the utilities for one of the choices freely, say for *tent.* To make things easy for us, we can choose all utilities for *tent* equal to 0, that is,

$$u(\textit{tent},\ \textit{rainy}) = 0,\ u(\textit{tent},\ \textit{stormy}) = 0 \text{ and } u(\textit{tent},\ \textit{calm}) = 0.$$

In Section 2.3.2 we have seen that if we start from an expected utility representation, we can also multiply all utilities by the same positive number and still be sure that we have an expected utility representation at the end. As such, the utility for *house* at state *rainy* can be chosen arbitrarily, as long as it is larger than 0, since you prefer *house* to *tent* at this state. Let us choose

$$u(\textit{house},\ \textit{rainy}) = 1.$$

Now, consider the state *calm.* What should the utility for *house* at this state be? To answer this question, concentrate on the line segment of beliefs that only assign positive probability to the states *rainy* and *calm.* This line segment goes from the state *rainy* to the state *calm.* Note that the belief $p_1 = (0.3, 0, 0.7)$ is on this line segment, and that you are indifferent between *house* and *tent* at this belief. That is, at the belief p_1, the expected utility difference between *house* and *tent* must be 0.

The expected utility difference between *house* and *tent,* when viewed as a function of the belief p on the line segment between *rain* and *calm,* can be vizualized by Figure 2.4.2. From the picture it can

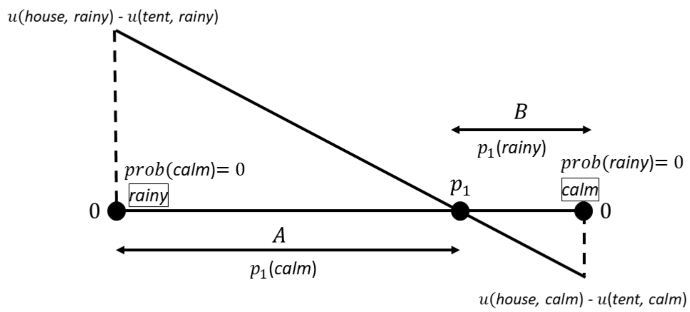

Figure 2.4.2 Expected utility difference on a line of beliefs

be seen that

$$\frac{u(house,\ rainy) - u(tent,\ rainy)}{u(tent,\ calm) - u(house,\ calm)} = \frac{A}{B}, \tag{2.4.1}$$

where A is the distance between *rainy* and p_1, and B is the distance between *calm* and p_1. Since the probability of *calm* at the state *rainy* is 0, the distance A is $p_1(calm) - 0 = p_1(calm)$. Similarly, the probability of *rainy* at the state *calm* is 0, and thus the distance B is $p_1(rainy) - 0 = p_1(rainy)$. Together with (2.4.1) we can therefore conclude that

$$\frac{u(house,\ rainy) - u(tent,\ rainy)}{u(tent,\ calm) - u(house,\ calm)} = \frac{A}{B} = \frac{p_1(calm)}{p_1(rainy)}. \tag{2.4.2}$$

This formula enables us to compute $u(house,\ calm)$, since we already know $u(tent,\ rainy)$, $u(tent,\ calm)$ and $u(house,\ rainy)$. Indeed, since $u(tent,\ rainy) = u(tent,\ calm) = 0$ and $u(house,\ rainy) = 1$, it follows from (2.4.2) that

$$\frac{1}{-u(house,\ calm)} = \frac{p_1(calm)}{p_1(rainy)} = \frac{0.7}{0.3} = \frac{7}{3},$$

which yields

$$u(house,\ calm) = -3/7.$$

In fact, the formula (2.4.2) above can be generalized for any two states, any two choices, and any belief on the line between the two states where the DM is indifferent between the two choices. We call this the *utility difference property.*

Utility difference property. *Consider a conditional preference relation $\succsim$ with an expected utility representation u. Take two choices a, b and two states x, y such that the DM prefers a to b at x and prefers b to a at y. Suppose that the DM is indifferent between a and b at the belief p on the line segment between $[x]$ and $[y]$. Then,*

$$\frac{u(a,x) - u(b,x)}{u(b,y) - u(a,y)} = \frac{p(y)}{p(x)}.$$

We can use the utility difference property to finally compute the utility for *house* at the state *stormy.* Let $[s]$ be the belief that assigns probability 1 to state *stormy.* Consider the line segment

u	*rainy*	*stormy*	*calm*
tent	0	0	0
house	1	17/7	−3/7

v	*rainy*	*stormy*	*calm*
tent	3	0	3
house	10	17	0

Table 2.4.1 Two expected utility representations for the conditional preference relation in Figure 2.4.1

of beliefs between $[s]$ and $[c]$, and the belief $p_2 = (0, 0.15, 0.85)$ on this line segment where you are indifferent between *house* and *tent.* Then, it follows by the utility difference property that

$$\frac{u(\textit{house, stormy}) - u(\textit{tent, stormy})}{u(\textit{tent, calm}) - u(\textit{house, calm})} = \frac{p_2(\textit{calm})}{p_2(\textit{stormy})}.$$

Since $u(\textit{tent, stormy}) = 0$, $u(\textit{tent, calm}) = 0$, $u(\textit{house, calm}) = -3/7$, $p_2(\textit{calm}) = 0.85$ and $p_2(\textit{stormy}) = 0.15$, we obtain

$$\frac{u(\textit{house, stormy})}{3/7} = \frac{0.85}{0.15} = \frac{17}{3}.$$

Thus,

$$u(\textit{house, stormy}) = 17/7.$$

Hence, we obtain the utility function u in the left-hand panel of Table 2.4.1.

It may be verified that this utility function *represents* the conditional preference relation from Figure 2.4.1. To see this, consider the utility function v in the right-hand panel of Table 2.4.1, which is based on the utility function from Table 2.3.1. We have already seen in the previous section that v represents the conditional preference relation from Figure 2.4.1.

Note that at each of the three states, the utility difference in v is exactly 7 times as large as in u. Consequently, at every belief the *expected* utility difference induced by v is exactly 7 times as large compared to u. In particular, both utility functions will induce the same preferences between *tent* and *house* at each of the beliefs. Since we have already seen that v represents the conditional preference relation in Figure 2.4.1, the same holds for utility function u.

Question 2.4.1 *Consider the conditional preference relation in Figure 2.3.1, and concentrate on the preferences between garden and tent. In a similar way as above, derive the unique utility function u that represents this conditional preference relation between garden and tent, such that the utilities for tent are always* 0, *and* $u(\textit{garden, calm}) = 1$.

The method we used above to calculate an expected utility representation for the conditional preference relation in Figure 2.4.1 can be generalized. The resulting procedure, called *utility design procedure for two choices with reversals,* can be applied to conditional preference relations on two choices, a and b, where there are *preference reversals* between a and b. Recall from Definition 2.2.5 that there are preference reversals between a and b if there is some belief p with $a \succ_p b$ and some belief q with $b \succ_q a$. If, in addition, there is an expected utility representation for the conditional preference relation, then it can be shown that there must be a state x with $a \succ_{[x]} b$ and another state y with $b \succ_{[y]} a$. Here, $[x]$ denotes the belief that assigns probability 1 to the state x, and similarly for $[y]$.

Definition 2.4.1 (Utility design for two choices with reversals) *Consider a conditional preference relation $\succsim$ with two choices, a and b, such that there are preference reversals between a and b.*

To start, define $u(a, s)$ arbitrarily for every state s, select two states, x and y, such that $a \succ_{[x]} b$ and

$b \succ_{[y]} a$, and set $u(b,y) = \alpha$ for some arbitrary number $\alpha > u(a,y)$.

Select a belief p on the line between $[x]$ and $[y]$ with $a \sim_p b$, and determine $u(b,x)$ by the utility difference property

$$\frac{u(a,x) - u(b,x)}{u(b,y) - u(a,y)} = \frac{p(y)}{p(x)}.$$

For every state $s \neq x, y$ with $b \succ_{[s]} a$, select a belief p on the line between $[x]$ and $[s]$ with $a \sim_p b$, and determine $u(b,s)$ by the utility difference property

$$\frac{u(b,s) - u(a,s)}{u(a,x) - u(b,x)} = \frac{p(x)}{p(s)}.$$

For every state $s \neq x, y$ with $a \succ_{[s]} b$, select a belief p on the line between $[y]$ and $[s]$ with $a \sim_p b$, and determine $u(b,s)$ by the utility difference property

$$\frac{u(a,s) - u(b,s)}{u(b,y) - u(a,y)} = \frac{p(y)}{p(s)}.$$

For every state $s \neq x, y$ with $a \sim_{[s]} b$, set $u(b,s) := u(a,s)$.

The procedure not only enables us to compute an expected utility representation if there is some, but also helps to determine whether the conditional preference relation has an expected utility representation *at all*. Indeed, suppose that we start from a conditional preference relation $\succsim$, use the procedure above and get stuck because there is no belief p on the line between $[x]$ and $[y]$ with $a \sim_p b$. Then, we can conclude that there is no expected utility representation. The intuition is clear: If there would be an expected utility representation u, then $u(a,x) > u(b,x)$ and $u(b,y) > u(a,y)$, and hence there would be a belief p on the line between $[x]$ and $[y]$ with $u(a,p) = u(b,p)$, which would imply that $a \sim_p b$. Similarly, if at some of the further steps we conclude that there is no belief on the line between $[x]$ and $[s]$, or on the line between $[y]$ and $[s]$, with $a \sim_p b$, then this means there is no expected utility representation.

Or suppose we are able to run the procedure until the end, find a utility function u, but conclude that u does not represent the conditional preference relation $\succsim$. Also in this case we may safely conclude that there is no expected utility representation at all. As an illustration, consider the conditional preference relation $\succsim$ in Figure 2.4.3 for the "Birthday party" with choices *tent* and *house.*

It is easily seen that this conditional preference relation can have no expected utility representation. Indeed, suppose it would have an expected utility representation u. Then, $u(\textit{house}, p_1) = u(\textit{tent}, p_1)$ and $u(\textit{house}, p_2) = u(\textit{tent}, p_2)$, which would imply that also $u(\textit{house}, p) = u(\textit{tent}, p)$ for all the beliefs p on the line between p_1 and p_2. That is, you should be indifferent between *house* and *tent* for all beliefs on the line between p_1 and p_2. However, from the figure it is clear that you prefer *tent* to *house* for all beliefs on the line between p_1 and p_2, except for the beliefs p_1 and p_2 themselves. This means there can be no expected utility representation.

Now suppose you would run the utility design procedure to compute a utility function u. As the procedure is completely based on the beliefs p_1 and p_2 this would lead to the utility function u we computed above, represented in the left-hand panel of Table 2.4.1. Consider the belief $p = (0.15, 0.075, 0.775)$ exactly halfway between p_1 and p_2. Then, it may be verified that $u(\textit{house}, p) = u(\textit{tent}, p)$, whereas the conditional preference relation indicates that $t \succ_p h$. As such, the utility function u so computed does not represent the conditional preference relation $\succsim$. Based on this, we conclude that $\succsim$ has no expected utility representation at all.

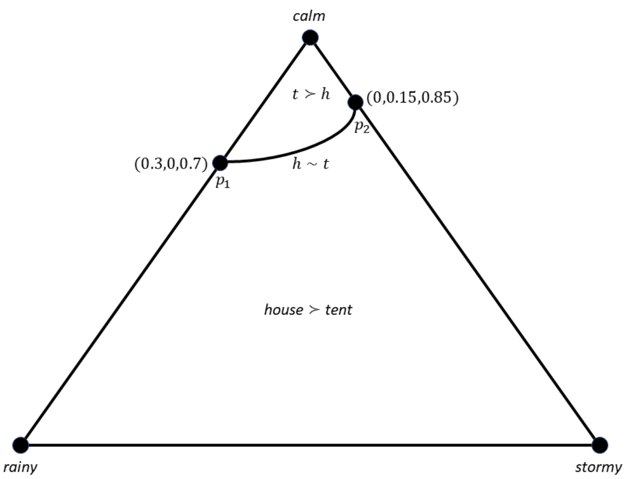

Figure 2.4.3 When there is no expected utility representation

To conclude, we illustrate the utility design procedure by an example with four states. The story is as follows: Last year you had a terrible camping experience, and for that reason you have discarded the *tent* as a possible location for your birthday party. Hence, you must decide between having the party in your *house* and having it in the *garden.* You realized that the category *calm* was a bit too vague to describe the weather, because *calm* weather could mean *sunny* or *cloudy,* and this may make a difference for your optimal choice. You therefore replace the state *calm* by two more descriptive states, *sunny* and *cloudy.*

Suppose that your conditional preference relation between *house* and *garden* is given by Figure 2.4.4. This is a three-dimensional picture. The four corner points of the pyramid represent the probability 1 beliefs, where you assign probability 1 to *rainy,* probability 1 to *cloudy,* probability 1 to *stormy*, and probability 1 to *sunny,* respectively. Similarly as for three states, every belief can be identified with a point in the pyramid. The closer a belief is to a certain cornerpoint, the higher the probability that the belief assigns to the state at that cornerpoint. For instance, the higher the probability that is being assigned to *rainy,* the closer the belief will be to the cornerpoint *rainy.* The six edges of the pyramid represent beliefs that assign positive probability to at most two states. As an example, the edge between *rainy* and *sunny* represents all beliefs that assign positive probability only to *rainy* and *sunny,* and assign probability 0 to the other two states.

The set of beliefs where you are different between *house* and *garden* is represented by the grey plane inside the pyramid, with four cornerpoints, including p_1, p_2 and p_3. For all beliefs above the plane you prefer *garden* whereas you prefer *house* for all beliefs below the plane. In particular, there are preference reversals between *house* and *garden.* We will now use the utility design procedure to find a utility function that represents this conditional preference relation.

We start by setting all utilities for *house* equal to 0, by fixing the states $x = rainy$ and $y = sunny$, and by setting $u(garden,\ sunny) = 1$. Then, we calculate $u(garden,\ rainy)$ by the utility difference property with respect to the belief $p_1 = 0.5\ rainy + 0.5\ sunny.$ That is,

$$\frac{u(house,\ rainy) - u(garden,\ rainy)}{u(garden,\ sunny) - u(house,\ sunny)} = \frac{p_1(sunny)}{p_1(rainy)} = \frac{0.5}{0.5} = 1.$$

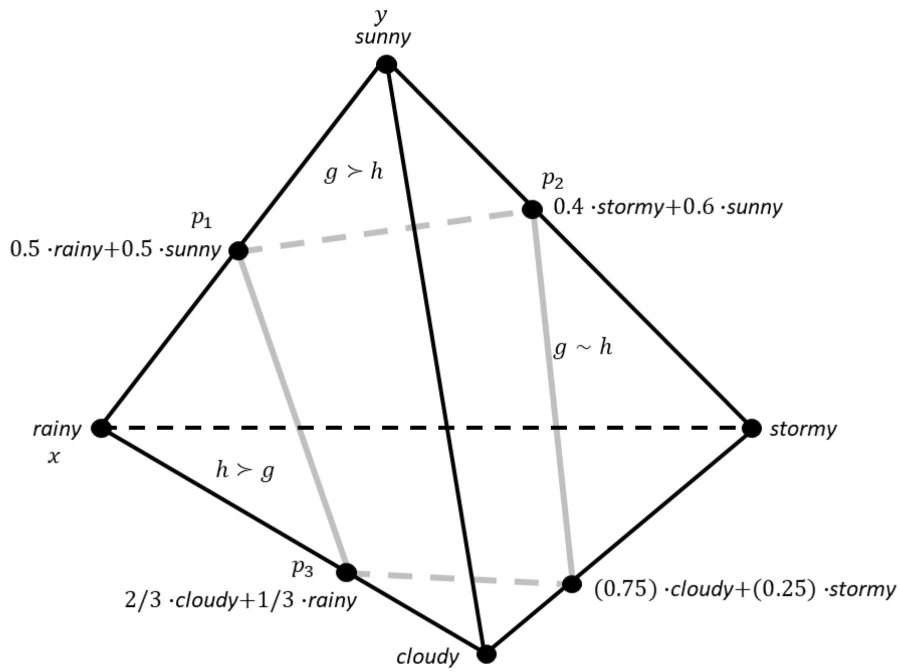

Figure 2.4.4 A conditional preference relation for four states of weather

Since $u(house,\ rainy) = u(house,\ sunny) = 0$ and $u(garden,\ sunny) = 1$, we get

$$\frac{-u(garden,\ rainy)}{1} = 1$$

and hence

$$u(garden,\ rainy) = -1.$$

Next, we compute $u(garden,\ stormy)$ by the utility difference property with respect to the belief $p_2 = 0.4$ *stormy* $+0.6$ *sunny.* That is,

$$\frac{u(house,\ stormy) - u(garden,\ stormy)}{u(garden,\ sunny) - u(house,\ sunny)} = \frac{p_2(sunny)}{p_2(stormy)} = \frac{0.6}{0.4} = 1.5.$$

As $u(house,\ stormy) = u(house,\ sunny) = 0$ and $u(garden,\ sunny) = 1$, we get

$$\frac{-u(garden,\ stormy)}{1} = 1.5,$$

which yields

$$u(garden,\ stormy) = -1.5.$$

Finally, we obtain $u(garden,\ cloudy)$ by the utility difference property with respect to the belief $p_3 = 2/3$ *cloudy* $+1/3$ *rainy.* That is,

$$\frac{u(garden,\ cloudy) - u(house,\ cloudy)}{u(house,\ rainy) - u(garden,\ rainy)} = \frac{p_3(rainy)}{p_3(cloudy)} = \frac{1/3}{2/3} = 0.5.$$

u	*sunny*	*rainy*	*stormy*	*cloudy*
house	0	0	0	0
garden	1	-1	-1.5	0.5

Table 2.4.2 An expected utility representations for the conditional preference relation in Figure 2.4.4

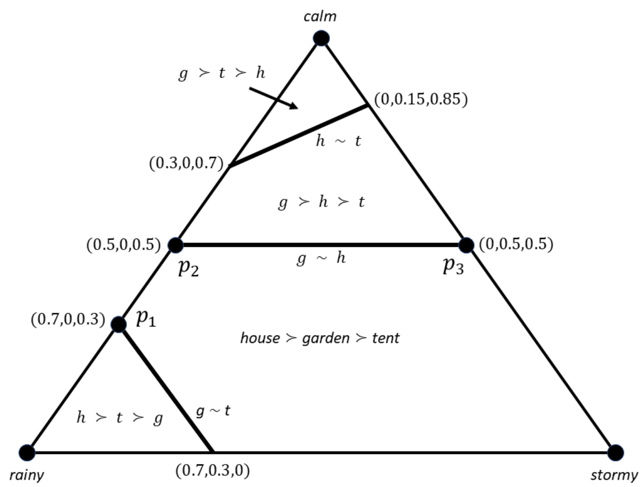

Figure 2.4.5 Utility design procedure for three choices

As $u(house,\ cloudy) = u(house,\ rainy) = 0$ and $u(garden,\ rainy) = -1$, we get

$$\frac{u(garden,\ cloudy)}{1} = 0.5,$$

which yields

$$u(garden,\ cloudy) = 0.5.$$

We thus obtain the utility function in Table 2.4.2. It may be verified that this utility function represents the conditional preference relation in Figure 2.4.4.

2.4.2 More Than Two Choices with Preference Reversals

We now move from two choices to three choices or more, but assume that there are preference reversals between every two choices. Similarly to what we have done for the case of two choices, we will present a *utility design procedure* for this case. To illustrate how the procedure works, let us return to the conditional preference relation in Figure 2.3.1, which we have reproduced in Figure 2.4.5. We have seen in Section 2.3 that it has an expected utility representation. But how would we compute these utilities, based on the data in Figure 2.4.5?

In the previous section we have already computed some utilities that represent the conditional preferences between *house* and *tent.* This was done based on the utility design procedure for two choices, and the resulting utilities have been reproduced in Table 2.4.3.

We next turn to the utilities for *garden.* How can these be derived? Note that the belief p_1 is on the line segment between *rainy* and *calm,* that the DM prefers *garden* to *tent* at the state *calm* and

u	*rainy*	*stormy*	*calm*
tent	0	0	0
house	1	17/7	−3/7

Table 2.4.3 Expected utility representation for the conditional preferences between *house and garden*

prefers *tent* to *garden* at the state *rainy.* By applying the utility difference property to the belief p_1 and the choices *garden* and *tent* we obtain

$$\frac{u(garden,\ calm) - u(tent,\ calm)}{u(tent,\ rainy) - u(garden,\ rainy)} = \frac{p_1(rainy)}{p_1(calm)}.$$

By substituting $p_1(rainy) = 0.7$, $p_1(calm) = 0.3$, $u(tent,\ calm) = 0$ and $u(tent,\ rainy) = 0$, it follows that

$$\frac{u(garden,\ calm)}{-u(garden,\ rainy)} = \frac{0.7}{0.3} = \frac{7}{3}.$$

By cross-multiplying we get

$$u(garden,\ calm) = -\frac{7}{3} \cdot u(garden,\ rainy). \tag{2.4.3}$$

Next, applying the utility difference property to the belief p_2 and the choices *garden* and *house* yields

$$\frac{u(garden,\ calm) - u(house,\ calm)}{u(house,\ rainy) - u(garden,\ rainy)} = \frac{p_2(rainy)}{p_2(calm)}.$$

By substituting $p_2(rainy) = 0.5$, $p_2(calm) = 0.5$, $u(house,\ calm) = -3/7$ and $u(house,\ rainy) = 1$, we get

$$\frac{u(garden,\ calm) + 3/7}{1 - u(garden,\ rainy)} = 1.$$

Thus,

$$u(garden,\ calm) = 4/7 - u(garden,\ rainy). \tag{2.4.4}$$

By combining (2.4.4) and (2.4.3) we obtain

$$-\frac{7}{3} \cdot u(garden,\ rainy) = 4/7 - u(garden,\ rainy),$$

which yields

$$\frac{4}{3} \cdot u(garden,\ rainy) = -\frac{4}{7},$$

and hence

$$u(garden,\ rainy) = -3/7.$$

If we substitute this into (2.4.4) we get

$$u(garden,\ calm) = 1.$$

Finally, if we apply the utility difference property to the belief p_3 and the choices *garden* and *house,* we obtain

$$\frac{u(garden,\ calm) - u(house,\ calm)}{u(house,\ stormy) - u(garden,\ stormy)} = \frac{p_3(stormy)}{p_3(calm)}.$$

	rainy	*stormy*	*calm*
house	1	17/7	−3/7
garden	−3/7	1	1
tent	0	0	0

	rainy	*stormy*	*calm*
house	10	17	0
garden	0	7	10
tent	3	0	3

Table 2.4.4 Expected utility representations for the conditional preference relation in Figure 2.4.5

Since $p_3(\textit{stormy}) = 0.5$, $p_3(\textit{calm}) = 0.5$, $u(\textit{garden, calm}) = 1$, $u(\textit{house, calm}) = -3/7$ and $u(\textit{house, stormy}) = 17/7$, it follows that

$$\frac{10/7}{17/7 - u(\textit{garden, stormy})} = 1.$$

This implies

$$u(\textit{garden, stormy}) = 1.$$

We thus obtain the utilities in the left-hand panel of Table 2.4.4.

It may be verified that these utilities represent the conditional preference relation at hand. The easiest way to see this is to compare these utilities to the ones in the right-hand panel of Table 2.4.4. We have seen in Section 2.3 that the right-hand utilities represent the conditional preference relation. Note that the right-hand utilities can be obtained from the left-hand utilities if we first multiply all of these utilities by 7, then add 3 to all utilities at the state *rainy,* and finally add 3 to all utilities at the state *calm.* As these transformation do not change the relative expected utility differences at any belief, it follows that the left-hand utilities, which we have obtained through the procedure above, will also represent the conditional preference relation.

The procedure we have used above can be generalized, and leads to the *utility design procedure.* It can be applied to any conditional preference relation with three or more choices where there are preference reversals for every pair of choices, and such that there is a belief where the DM is indifferent between some, but not all, choices. In the description of the procedure we denote, for every two choices a and b, by $P_{a\sim b}$ the set of beliefs where the DM is indifferent between a and b.

Definition 2.4.2 (Utility design for more than two choices with reversals) *Consider a conditional preference relation with more than two choices such that (i) there are preference reversals for every pair of choices, and (ii) there is a belief where the DM is indifferent between some, but not all, choices.*

Start by selecting three choices a, b, c such that the indifference sets $P_{a\sim c}$ and $P_{b\sim c}$ are different. Concentrate on the conditional preferences between a and b, and use the utility design procedure for two choices to find utilities for a and b.

To find the utilities for c, select first a belief p_1 where the DM is indifferent between c and a, but not between c and b.
If there are n states in total, number the states $s_1, s_2, ..., s_n$ such that $c \succ_{s_1} b$ and $b \succ_{s_2} c$, and find a belief p_2 on the line segment between $[s_1]$ and $[s_2]$ where $c \sim_{p_2} b$.
For every state s_k with $k \geq 3$ and $c \succ_{s_k} b$, find a belief p_k on the line between $[s_2]$ and $[s_k]$ where $c \sim_{p_k} b$.
For every state s_k with $k \geq 3$ and $b \succ_{s_k} c$, find a belief p_k on the line between $[s_1]$ and $[s_k]$ where $c \sim_{p_k} b$.

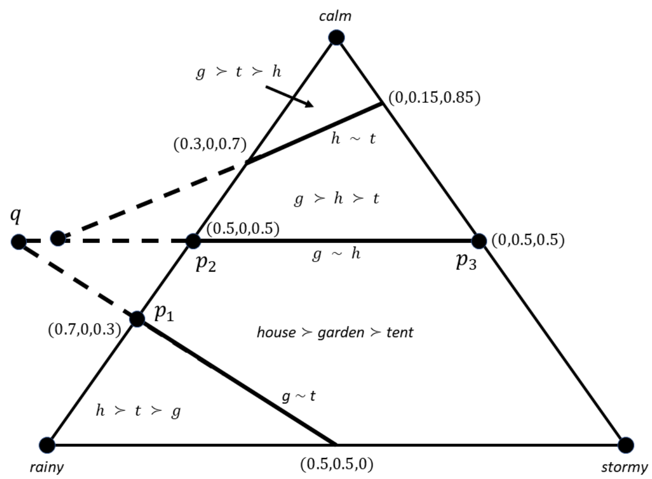

Figure 2.4.6 When there is no expected utility representation

For every state s_k with $k \geq 3$ and $c \sim_{s_k} b$, set $p_k = [s_k]$.
Then, by applying the utility difference property to the belief p_1 and the choices c and a, and by applying the utility difference property to the beliefs $p_2, ..., p_n$ and the choices c and b, we obtain a system with n equations and n variables, $u(c, s_1), ..., u(c, s_n)$. Find the unique values for $u(c, s_1), ..., u(c, s_n)$ that solve this system of equations. This yields the utilities for choice c

For every choice d different from a, b, c, there are two possible cases:

(a) if $P_{a \sim d}$ is different from $P_{b \sim d}$, then find the utilities for d as above, but now applied to the choices a, b and d instead of a, b and c;

(b) if $P_{a \sim d}$ is different from $P_{c \sim d}$, then find the utilities for d as above, but now applied to the choices a, c and d instead of a, b and c.

This procedure may seem a bit overwhelming, but things will become clear once we apply the procedure to some concrete examples. Let us first return to the procedure we have applied for the conditional preference relation in Figure 2.4.5. In terms of the utility design procedure above, we have set $a =$ *tent,* $b =$ *house* and $c =$ *garden.* Moreover, to compute the utilities for $c =$ *garden,* we have numbered the states by $s_1 =$ *calm,* $s_2 =$ *rainy* and $s_3 =$ *stormy.* Finally, we have chosen the beliefs p_1, p_2 and p_3 as indicated in Figure 2.4.5.

Similarly to what we have seen for two choices, also now the procedure can be used to verify whether the conditional preference relation at hand has an expected utility representation at all. It may happen, for instance, that we get stuck in the procedure because at some step we cannot find a belief p on the line between $[s_1]$ and $[s_k]$, or between $[s_2]$ and $[s_k]$, where $c \sim_p b$. In this case we may safely conclude that there is no expected utility representation for the conditional preference relation.

But it may also happen that we are able to fully implement the procedure, yet find out that the utility function u we obtain does not represent the conditional preference relation $\succsim$. Also in that case, we must conclude that $\succsim$ has no expected utility representation.

To illustrate this, consider the conditional preference relation $\succsim$ in Figure 2.4.6, for "The birthday party". Note that this conditional preference relation is almost the same as in Figure 2.4.5. However,

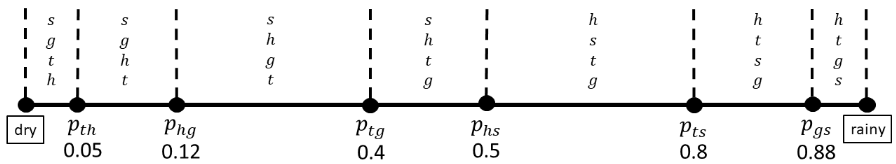

Figure 2.4.7 The birthday party with two states and four choices

the belief on the line between *rainy* and *stormy* where you are indifferent between *garden* and *tent* is now $(0.5, 0.5, 0)$ instead of $(0.7, 0.3, 0)$.

It turns out that this new conditional preference relation has no expected utility representation. To see this, suppose that there would be an expected utility representation u. Then, the expected utility for *garden* would be equal to the expected utility of *tent* for all beliefs on the line through $(0.7, 0, 0.3)$ and $(0.5, 0, 5, 0)$. We could even look at "signed" beliefs outside the belief triangle, where some "probabilities" may be negative, and conclude that at the signed belief q, which lies on the line through $(0.7, 0, 0.3)$ and $(0.5, 0, 5, 0)$, the "expected utility" for *garden* and *tent* must be the same.

Similarly, the expected utility for *garden* would be equal to the expected utility of *house* for all beliefs on the line through the beliefs $(0.5, 0, 0.5)$ and $(0, 0, 5, 0.5)$. Since the signed belief q lies on the line through $(0.5, 0, 0.5)$ and $(0, 0.5, 0.5)$, the "expected utility" for *garden* and *house* must be the same at q. As we have already seen that the "expected utility" for *garden* and *tent* must be the same at q, it follows that the "expected utility" of *tent* and *house* must be the same at q. However, q does not lie on the line through the beliefs $(0.3, 0, 0.7)$ and $(0, 0.15, 0.85)$ where the expected utility of *tent* and *house* are the same, which yields a contradiction. We thus conclude that the conditional preference relation does not have an expected utility representation.

This conclusion will also be reached if we run the utility design procedure above. If in the procedure we set $a =$ *tent,* $b =$ *house* and $c =$ *garden*, and number the states by $s_1 =$ *calm,* $s_2 =$ *rainy* and $s_3 =$ *stormy,* then the procedure will only depend on the beliefs where you are indifferent between *tent* and *house,* and the beliefs p_1, p_2 and p_3. Since all these beliefs are the same as in Figure 2.4.5, we find the same utility function u that we found for Figure 2.4.5, which is represented in the left-hand panel of Table 2.4.4. However, for this utility function you would not be indifferent between *garden* and *tent* at the belief $(0.5, 0.5, 0)$, and hence this utility function does not represent the conditional preference relation in Figure 2.4.6. On the basis of this insight, we may conclude that the conditional preference relation in Figure 2.4.6 does not have an expected utility representation.

To illustrate the procedure for the case of more than three choices, suppose that in the example "The birthday party" you replace the states *stormy* and *calm* by the "summary" state *dry.* Hence, for your decision where to organize the party you only distinguish between two states of weather, *rainy* and *dry.* At the same time, you thought about another possible location for your party, which is a nice *square* close to your house. Assume that your conditional preference relation is given by Figure 2.4.7.

This picture should be read as follows: The leftmost point represents the belief where you assign probability 1 to it being *dry,* whereas the rightmost point is the other extreme belief that assigns probability 1 to it being *rainy* today. Every belief can thus be identified with a point on the line between *dry* and *rainy.* The closer the point is to *rainy,* the higher the probability you assign to *rainy.* The six beliefs on the line, $p_{th}, p_{hg}, ..., p_{gs}$, represent the (unique) beliefs where you are indifferent between two of the choices. For instance, at the belief p_{th} you are indifferent between *tent* and *house,* and at the belief p_{gs} you are indifferent between *garden* and *square.* The numbers below these beliefs $p_{th}, p_{hg}, ..., p_{gs}$ indicate the probability assigned by that belief to the state *rainy.* Hence, the belief p_{th} assigns probability 0.05 to it being *rainy,* and therefore probability 0.95 to it being *dry.* Similarly for the other five beliefs.

These six beliefs partition the set of beliefs into seven different areas, and for each of these areas we have indicated your preference over the four choices by a vertical ranking. For instance, in the first area, where you assign a very high probability to *dry* weather, you prefer *square* to *garden, garden* to *tent,* and *tent* to *house.* In contrast, in the last area, where you assign a very high probability to *rainy* weather, your preference over the choices is reversed.

It turns out that this conditional preference relation has an expected utility representation. Moreover, there are preference reversals for all pairs of choices, and there is a belief where you are indifferent between some, but not all, choices. How would we compute an expected utility representation here, following the procedure above?

For the choices a, b and c we choose

$$a = \textit{ tent, } b = \textit{ house} \text{ and } c = \textit{ garden.}$$

We start by finding utilities for *tent* and *house,* using the utility design procedure for two choices. We may set all utilities for *tent* equal to 0, and choose $u(\textit{house, rainy}) = 1$. By applying the utility difference property to the belief $p_{th} = (0.05, 0.95)$, where the first number denotes the probability assigned to *rainy,* in combination with the choices *tent* and *house,* we get

$$\frac{u(\textit{house, rainy}) - u(\textit{tent, rainy})}{u(\textit{tent, dry}) - u(\textit{house, dry})} = \frac{p_{th}(\textit{dry})}{p_{th}(\textit{rainy})}.$$

If we subtitute $p_{th}(\textit{dry}) = 0.95$, $p_{th}(\textit{rainy}) = 0.05$, $u(\textit{house, rainy}) = 1$, $u(\textit{tent, rainy}) = 0$ and $u(\textit{tent, dry}) = 0$, we get

$$\frac{1 - 0}{0 - u(\textit{house, dry})} = \frac{0.95}{0.05} = 19,$$

and hence

$$u(\textit{house, dry}) = -1/19.$$

To compute the utilities for choice $c = \textit{garden}$, we use the beliefs $p_1 = p_{tg} = (0.4, 0.6)$ and $p_2 = p_{hg} = (0.12, 0.88)$. If we apply the utility difference property to the belief p_1 for the choices *garden* and *tent,* we obtain

$$\frac{u(\textit{tent, rainy}) - u(\textit{garden, rainy})}{u(\textit{garden, dry}) - u(\textit{tent, dry})} = \frac{p_{tg}(\textit{dry})}{p_{tg}(\textit{rainy})}.$$

By filling in $p_{tg}(\textit{dry}) = 0.6$, $p_{tg}(\textit{rainy}) = 0.4$, $u(\textit{tent, rainy}) = 0$ and $u(\textit{tent, dry}) = 0$, we get

$$\frac{0 - u(\textit{garden, rainy})}{u(\textit{garden, dry}) - 0} = \frac{0.6}{0.4} = \frac{3}{2}.$$

Cross-multiplication then yields

$$u(\textit{garden, rainy}) = -3/2 \cdot u(\textit{garden, dry}). \tag{2.4.5}$$

Applying the utility difference property to the belief p_2 for the choices *garden* and *house* yields

$$\frac{u(\textit{house, rainy}) - u(\textit{garden, rainy})}{u(\textit{garden, dry}) - u(\textit{house, dry})} = \frac{p_{hg}(\textit{dry})}{p_{hg}(\textit{rainy})}.$$

Substituting $p_{hg}(dry) = 0.88$, $p_{hg}(rainy) = 0.12$, $u(house,\ rainy) = 1$ and $u(house,\ dry) = -1/19$ then leads to

$$\frac{1 - u(garden,\ rainy)}{u(garden,\ dry) + 1/19} = \frac{0.88}{0.12} = \frac{22}{3}.$$

By cross-multiplication, we get

$$3 - 3 \cdot u(garden,\ rainy) = 22 \cdot u(garden,\ dry) + 22/19$$

which yields

$$u(garden,\ rainy) = 35/57 - 22/3 \cdot u(garden,\ dry).$$

If we substitute this into (2.4.5), we get

$$35/57 - 22/3 \cdot u(garden,\ dry) = -3/2 \cdot u(garden,\ dry).$$

Solving for $u(garden,\ dry)$ then yields

$$u(garden,\ dry) = 2/19.$$

By (2.4.5) we then get

$$u(garden,\ rainy) = -3/19.$$

Finally, for computing the utilities of *square* we use the beliefs $p_1 = p_{ts} = (0.8, 0.2)$ and $p_2 = p_{hs} = (0.5, 0.5)$. If we apply the utility difference property to the belief p_1 and the choices *square* and *tent,* we get

$$\frac{u(tent,\ rainy) - u(square,\ rainy)}{u(square,\ dry) - u(tent,\ dry)} = \frac{p_{ts}(dry)}{p_{ts}(rainy)}.$$

If we substitute $p_{ts}(dry) = 0.2$, $p_{ts}(rainy) = 0.8$, $u(tent,\ rainy) = 0$ and $u(tent,\ dry) = 0$, we obtain

$$\frac{0 - u(square,\ rainy)}{u(square,\ dry) - 0} = \frac{0.2}{0.8} = \frac{1}{4}.$$

Cross-multiplication then leads to

$$u(square,\ dry) = -4 \cdot u(square,\ rainy). \tag{2.4.6}$$

By applying the utility difference property to the belief p_2 and the choices *square* and *house,* we get

$$\frac{u(house,\ rainy) - u(square,\ rainy)}{u(square,\ dry) - u(house,\ dry)} = \frac{p_{hs}(dry)}{p_{hs}(rainy)}.$$

If we substitute $p_{hs}(dry) = 0.5$, $p_{hs}(rainy) = 0.5$, $u(house,\ rainy) = 1$ and $u(house,\ dry) = -1/19$, we obtain

$$\frac{1 - u(square,\ rainy)}{u(square,\ dry) + 1/19} = \frac{0.5}{0.5} = 1.$$

Hence,

$$u(square,\ dry) = 18/19 - u(square,\ rainy).$$

If we substitute this into (2.4.6), we get

$$18/19 - u(square,\ rainy) = -4 \cdot u(square,\ rainy),$$

	dry	*rainy*
tent	0	0
house	$-1/19$	1
garden	$2/19$	$-3/19$
square	$24/19$	$-6/19$

	dry	*rainy*
tent	0	0
house	-1	19
garden	2	-3
square	24	-6

Table 2.4.5 Expected utility representations for the conditional preference relation in Figure 2.4.7

which yields

$$u(\textit{square, rainy}) = -6/19.$$

Together with (2.4.6), we can conclude that

$$u(\textit{square, dry}) = 24/19.$$

We thus obtain the utility matrix in the left-hand panel of Table 2.4.5. It may be verified that this utility matrix indeed represents the conditional preference relation at hand. If we multiply each of the utilities by 19, we obtain the "easier" utility representation in the right-hand panel of Table 2.4.5.

Suppose now that we slightly change the conditional preference relation in Figure 2.4.7, by changing $p_{gs} = (0.88, 0.12)$ into $p'_{gs} = (0.9, 0.1)$, while leaving everything else unchanged. We will see that the new conditional preference relation does not have an expected utility representation. To show this, we can run the utility design procedure for the new conditional preference relation, using the same steps as we did for the "original" conditional preference relation in Figure 2.4.7. Note that the steps in the procedure only depend on the beliefs $p_{th}, p_{hg}, p_{tg}, p_{hs}$ and p_{ts}, which are the same as before. For that reason, we obtain the same utility function u as for the "original" conditional preference relation, depicted in the left-hand panel of Table 2.4.5. However, for this utility function you would not be indifferent between *garden* and *square* at the belief $p'_{gs} = (0.9, 0.1)$, and therefore the utility function we constructed does not represent the conditional preference relation. As such, we conclude that the new conditional preference relation does not have an expected utility representation.

2.4.3 *General Case

We finally turn to the general case, where there may be no preference reversals for some pairs of choices. It will be shown, by means of an example, how we can compute an expected utility representation for such cases.

Let us return to the birthday example which we know so well by now. Suppose you consider four possible locations for your party: *house, garden, tent* and *square.* Assume that you alway prefer *house* to *tent,* since you had a terrible camping experience last year. That is, your choice *house* strictly dominates your choice *tent.* Moreover, as the *square* is really far from your house, you always prefer *garden* to *square.* In other words, *garden* strictly dominates *square.* Indeed, both locations are in the open air, but the *garden* is much closer. Suppose that the corresponding conditional preference relation is given by Figure 2.4.8.

Let us first focus on the conditional preference relation $\succsim$, which is given by the objects inside the belief triangle. For the sake of clarity of the picture, we have only explicitly specified the preference relation over the four choices for one of the areas in the triangle. Otherwise, the picture would become too crowded. The preference relations for each of the other areas can be logically deduced from this single preference relation, by using the sets of indifference beliefs. For instance, in the upper area of

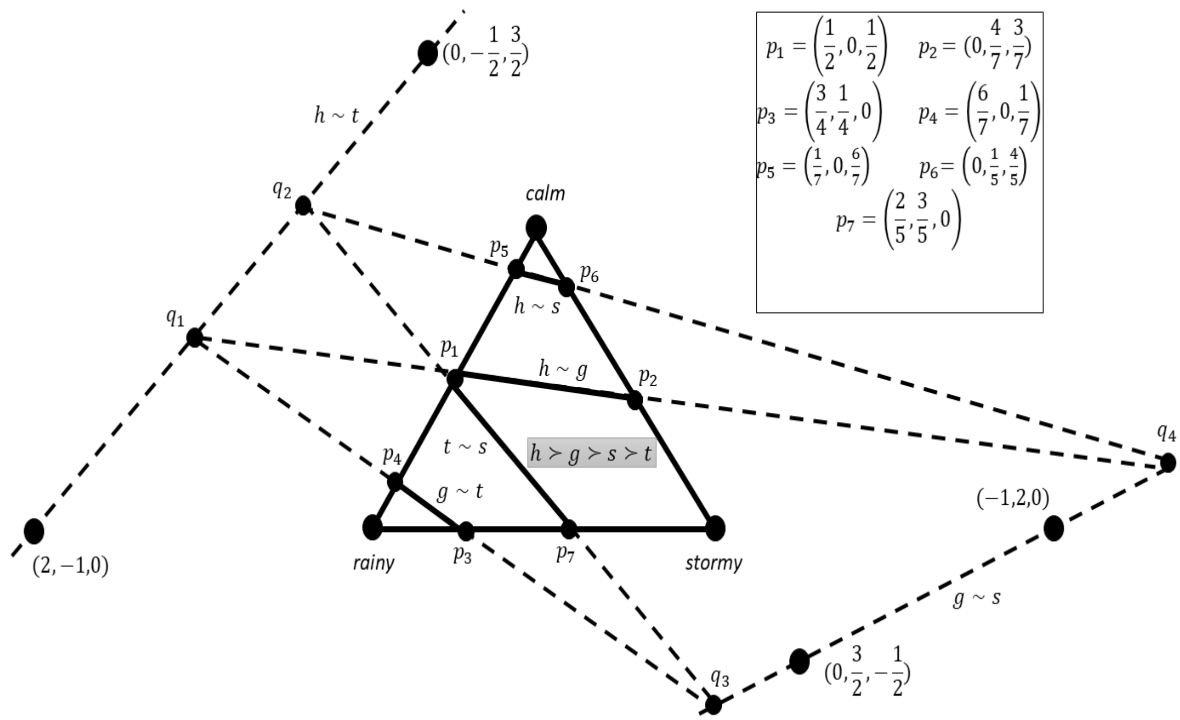

Figure 2.4.8 Utility design procedure when there are strictly dominated choices

the triangle we would have $g \succ s \succ h \succ t$. The reason is that if we start from the lower right area, where we know that $h \succ g \succ s \succ t$, and move to the upper area, then we cross the indifference set between h and g, and the indifference set between h and s. Thus, we first reverse the preferences between h and g, and afterwards reverse the preferences between h and s, which results in the ranking $g \succ s \succ h \succ t$.

Note that inside the belief triangle, there are no beliefs where you are indifferent between *house* and *tent,* or between *garden* and *square.* As such, the preferences between *house* and *tent,* and between *garden* and *square,* never get reversed in the belief triangle. We thus conclude that you will always prefer *house* to *tent,* and *garden* to *square,* for every belief. Or, in other words, your choice *house* strictly dominates *tent*, and your choice *garden* strictly dominates *square,* as already specified by the story above.

We now move to the objects outside the belief triangle. A point outside the belief triangle is called a *signed belief*, since some of the entries may be negative. For instance, the point $(2, -1, 0)$ in the picture is a signed belief that assigns value 2 to state *rainy,* negative value -1 to state *stormy,* and value 0 to state *calm.* Note that these values cannot be interpreted as probabilities, since some may be negative, or larger than 1. However, the sum of the values is still 1, like in a belief.

Definition 2.4.3 (Signed belief) *For a given set of states S, a* ***signed belief*** *q assigns to every state s a (possibly negative) number $q(s)$ such that $\sum_{s \in S} q(s) = 1$.*

Other signed beliefs in the picture are, for instance, q_1, q_2 and q_3. Note that every belief is a signed belief, but not *vice versa.* Recall that a conditional preference relation $\succsim$ specifies for every belief p a preference relation $\succsim_p$ over the choices. Similarly, we may define a *signed* conditional preference relation as an object that specifies a preference relation over choices for every *signed* belief.

Definition 2.4.4 (Signed conditional preference relation) *A* ***signed conditional preference relation*** $\succsim^*$ *assigns to every signed belief* q *a preference relation* $\succsim^*_q$ *over the choices.*

Every conditional preference relation $\succsim$ can be *extended* to a *signed* conditional preference relation $\succsim^*$ by additionally specifying a preference relation over choices for every signed belief that is not a belief. In Figure 2.4.8 we have depicted, by means of the dashed lines, a signed conditional preference relation $\succsim^*$ that extends $\succsim$. We have only specified the sets of signed beliefs where you are "indifferent" between a pair of choices, but by the same argument as above you can deduce the preference relation over the four locations for every area of signed beliefs.

Instead of computing an expected utility representation for the conditional preference relation $\succsim$, we will construct an expected utility representation u for the *signed* conditional preference relation $\succsim^*$ that extends it. By this, we mean that for every signed belief q, and every two choices a and b, we have that $a \succsim^*_q b$ precisely when $u(a,q) \geq u(b,q)$. Here,

$$u(a,q) := \sum_{s \in S} q(s) \cdot u(a,s)$$

denotes the "expected utility" induced by the choice a at the signed belief q, and similarly for $u(b,q)$.

To compute an expected utility representation u for the signed conditional preference relation $\succsim^*$ in Figure 2.4.8 we will use the same techniques as in Section 2.4.2, which applied to scenarios with preference reversals. We start by assigning some arbitrary utilities to the choice *house.* For instance, $u(\textit{house, rainy}) = 6$, $u(\textit{house, stormy}) = 6$ and $u(\textit{house, calm}) = 6$.

We now move to the utilities for *garden.* At the state *rainy,* we choose some arbitrary utility smaller than 6, because you prefer *house* to *garden* when it is *rainy.* For instance, $u(\textit{garden, rainy}) = 2$.

To compute $u(\textit{garden, calm})$, we use the utility difference property at the belief $p_1 = (1/2, 0, 1/2)$, where you are indifferent between *garden* and *house.* We obtain

$$\frac{u(\textit{garden, calm}) - u(\textit{house, calm})}{u(\textit{house, rainy}) - u(\textit{garden, rainy})} = \frac{p_1(\textit{rainy})}{p_1(\textit{calm})} = \frac{1/2}{1/2} = 1.$$

As $u(\textit{house, calm}) = 6$, $u(\textit{house, rainy}) = 6$ and $u(\textit{garden, rainy}) = 2$, we conclude that $u(\textit{garden, calm}) = 10$.

We then apply the utility difference property at the belief $p_2 = (0, 4/7, 3/7)$, where you are indifferent between *garden* and *house,* to compute $u(\textit{garden, stormy})$. This yields

$$\frac{u(\textit{garden, calm}) - u(\textit{house, calm})}{u(\textit{house, stormy}) - u(\textit{garden, stormy})} = \frac{p_2(\textit{stormy})}{p_2(\textit{calm})} = \frac{4/7}{3/7} = 4/3.$$

Since $u(\textit{garden, calm}) = 10$, $u(\textit{house, calm}) = 6$ and $u(\textit{house, stormy}) = 6$, it follows that $u(\textit{garden, stormy}) = 3$.

We next move to the utilities for *tent.* Note that the signed belief $(2, -1, 0)$ lies on the line through *rainy* and *stormy.* By applying the utility difference property to the signed beliefs $r_1 = (2, -1, 0)$, where you are "indifferent" between *tent* and *house,* we get

$$\frac{u(\textit{house, rainy}) - u(\textit{tent, rainy})}{u(\textit{tent, stormy}) - u(\textit{house, stormy})} = \frac{r_1(\textit{stormy})}{r_1(\textit{rainy})} = \frac{-1}{2} = -1/2.$$

As $u(\textit{house, rainy}) = 6$ and $u(\textit{house, stormy}) = 6$, we conclude that

$$u(\textit{tent, stormy}) = 2 \cdot u(\textit{tent, rainy}) - 6. \tag{2.4.7}$$

By applying the utility difference property to the belief $p_3 = (3/4, 1/4, 0)$, where you are indifferent between *tent* and *garden,* we get

$$\frac{u(\textit{garden, rainy}) - u(\textit{tent, rainy})}{u(\textit{tent, stormy}) - u(\textit{garden, stormy})} = \frac{p_3(\textit{stormy})}{p_3(\textit{rainy})} = \frac{1/4}{3/4} = \frac{1}{3}.$$

Since $u(\textit{garden, rainy}) = 2$ and $u(\textit{garden, stormy}) = 3$, this leads to

$$u(\textit{tent, stormy}) = 9 - 3 \cdot u(\textit{tent, rainy}). \tag{2.4.8}$$

By combining (2.4.7) and (2.4.8) we obtain

$$2 \cdot u(\textit{tent, rainy}) - 6 = 9 - 3 \cdot u(\textit{tent, rainy})$$

and hence $u(\textit{tent, rainy}) = 3$. From (2.4.7) it then follows that $u(\textit{tent, stormy}) = 0$.

Note that the signed belief $(0, -1/2, 3/2)$ is on the line through *stormy* and *calm.* To compute $u(\textit{tent, calm})$, we apply the utility difference property to the signed belief $r_2 = (0, -1/2, 3/2)$, where you are "indifferent" between *tent* and *house.* This yields

$$\frac{u(\textit{house, stormy}) - u(\textit{tent, stormy})}{u(\textit{tent, calm}) - u(\textit{house, calm})} = \frac{r_2(\textit{calm})}{r_2(\textit{stormy})} = \frac{3/2}{-1/2} = -3.$$

As $u(\textit{house, stormy}) = 6$, $u(\textit{tent, stormy}) = 0$ and $u(\textit{house, calm}) = 6$, it follows that $u(\textit{tent, calm}) = 4$.

We finally compute the utilities for *square.* Note that the signed belief $r_3 = (-1, 2, 0)$ is on the line through *rainy* and *stormy.* If we apply the utility difference property to the signed belief r_3, where you are "indifferent" between *square* and *garden,* we obtain

$$\frac{u(\textit{garden, rainy}) - u(\textit{square, rainy})}{u(\textit{square, stormy}) - u(\textit{garden, stormy})} = \frac{r_3(\textit{stormy})}{r_3(\textit{rainy})} = \frac{2}{-1} = -2.$$

Since $u(\textit{garden, rainy}) = 2$ and $u(\textit{garden, stormy}) = 3$, it follows that

$$u(\textit{square, rainy}) = 2 \cdot u(\textit{square, stormy}) - 4. \tag{2.4.9}$$

If we apply the utility difference property to the belief $p_7 = (2/5, 3/5, 0)$, where you are indifferent between *square* and *tent,* we get

$$\frac{u(\textit{tent, rainy}) - u(\textit{square, rainy})}{u(\textit{square, stormy}) - u(\textit{tent, stormy})} = \frac{p_7(\textit{stormy})}{p_7(\textit{rainy})} = \frac{3/5}{2/5} = 3/2.$$

As $u(\textit{tent, rainy}) = 3$ and $u(\textit{tent, stormy}) = 0$, we conclude that

$$u(\textit{square, rainy}) = 3 - (3/2) \cdot u(\textit{square, stormy}). \tag{2.4.10}$$

By combining (2.4.9) and (2.4.10) we obtain

$$2 \cdot u(\textit{square, stormy}) - 4 = 3 - (3/2) \cdot u(\textit{square, stormy}),$$

and hence $u(\textit{square, stormy}) = 2$. By substituting this into (2.4.9) we get that $u(\textit{square, rainy}) = 0$.

	rainy	*stormy*	*calm*
house	6	6	6
garden	2	3	10
tent	3	0	4
square	0	2	7

Table 2.4.6 Expected utility representation for the conditional preference relation in Figure 2.4.8

Finally, to compute $u(\textit{square, calm})$ we apply the utility difference property to the belief $p_1 = (1/2, 0, 1/2)$, where you are indifferent between *square* and *tent*. This yields

$$\frac{u(\textit{tent, rainy}) - u(\textit{square, rainy})}{u(\textit{square, calm}) - u(\textit{tent, calm})} = \frac{p_1(\textit{calm})}{p_1(\textit{rainy})} = \frac{1/2}{1/2} = 1.$$

Since $u(\textit{tent, rainy}) = 3$, $u(\textit{square, rainy}) = 0$ and $u(\textit{tent, calm}) = 4$, it follows that $u(\textit{square, calm}) = 7$.

We have thus computed all utilities. The resulting utility function is summarized by Table 2.4.6. It may be verified that this utility function indeed represents the signed conditional preference relation $\succsim^*$ in Figure 2.4.8. In particular, it represents the conditional preference relation $\succsim$.

2.5 Unique Relative Preference Intensities

Consider a conditional preference relation $\succsim$ that has an expected utility representation. In Section 2.3.2 we have seen that, if the DM is not always indifferent between each of his choices, then there are at least $|S| + 1$ degrees of freedom for selecting an expected utility representation, where $|S|$ is the number of states. One degree of freedom arises because we can multiply all utilities in the matrix by the same positive number $\alpha > 0$ and still obtain an expected utility representation. The other $|S|$ degrees of freedom arise because for every state s, we can always add the same number v_s to all the utilities at state s, and be sure that we will still have an expected utility representation.

Suppose now that there are preference reversals for every pair of choices, and that there is a belief where the DM is indifferent between some, but not all, choices (provided there are more than two choices). In that case, there will be *exactly* $|S| + 1$ degrees of freedom, and not more. The reason is as follows.

Note under the conditions above we can apply the utility design procedure for two choices, or more than two choices, with preference reversals to find an expected utility representation. In fact, *every* expected utility representation can be obtained through this utility design procedure. But how many degrees of freedom do we have in the procedure? First, we can choose the utilities for a at the $|S|$ different states freely, which gives us $|S|$ degrees of freedom already. Secondly, for one state s^* where $u(b, s^*) > u(a, s^*)$, we can choose the utility $u(b, s^*)$ freely, as long as $u(b, s^*) > u(a, s^*)$. This gives us the $|S| + 1$-th degree of freedom.

However, in the remainder of the procedure, where we find the utilities for b at all states other than s^*, and find the utilities for all other choices at every state, each of these remaining utilities will be unique determined. As such, the utility design procedure gives us exactly $|S| + 1$ degrees of freedom, and not more.

Since every expected utility representation can be found through this procedure, we conclude that there are precisely $|S|+1$ degrees of freedom for choosing the expected utility representation, assuming that the prerequisites for applying the procedure are met. That is, if we fix an expected utility representation u, then *every* other expected utility representation v can be found if we first multiply all utilities in u by a fixed number $\alpha > 0$, and subsequently add for every state s a fixed number v_s to all the utilities at state s.

But then, for all states, the *utility differences* between any pair of choices will be the same under the representations u and v, up to a positive multiplicative constant. Indeed, suppose that

$$v(a,s) = \alpha \cdot u(a,s) + v_s$$

for every choice a and every state s, where $\alpha > 0$. That is, we first multiply all utilities in u by α, and then add for every state s the number v_s to all utilities at s. Then, for every state s, and every two choices a and b, we have that

$$v(a,s) - v(b,s) = (\alpha \cdot u(a,s) + v_s) - (\alpha \cdot u(b,s) + v_s) = \alpha \cdot (u(a,s) - u(b,s)).$$

Thus, at every state the utility differences under u and v will be the same, up to a positive multiplicative constant α. In other words, the *relative* utility differences are *unique* at every state.

Recall from Section 2.3.3 that the relative utility differences at a given belief represent the relative *preference intensities* between the various choices at that belief. The insight above thus tells us that under the prerequisites mentioned above, the relative preference intensities between the choices will be unique. This is summarized by the following result.

Theorem 2.5.1 (Unique relative preference intensities) *Consider a conditional preference relation such that (i) it has an expected utility representation, (ii) there are preference reversals for all pairs of choices, and (iii) if there are more than two choices, there is a belief where the DM is indifferent between some, but not all, choices.*

Then, for every two expected utility representations u and v there is a positive number $\alpha > 0$ such that

$$v(a,s) - v(b,s) = \alpha \cdot (u(a,s) - u(b,s))$$

for every two choices a and b, and for every state s.

Consider, for instance, the two expected utility representations in Table 2.4.5 and call them u and v. Then, the utility differences in v are always 19 times as high as the utility differences in u. That is, $\alpha = 19$ in this case.

But what happens if the prerequisites in Theorem 2.5.1 are not satisfied? In this case, as we will see, there may be more than $|S|+1$ degrees of freedom for the expected utility representation, and the relative preference intensities may no longer be unique.

Consider, for instance, the conditional preference relation in Figure 2.5.1, with three choices a, b and c, and two states x and y. Note that the prerequisites for Theorem 2.5.1 are not met, since there is no belief where the DM is indifferent between some, but not all, beliefs. Indeed, the only belief where the DM is indifferent between two choices is $0.5[x] + 0.5[y]$, and at this belief he is indifferent between *all* three choices.

It may be verified that the two utility functions u and v in Table 2.5.1 both represent this conditional preference relation. Note that under the utility function u, the expected utility difference between a and b at state x is 1 whereas the expected utility difference between b and c at this state

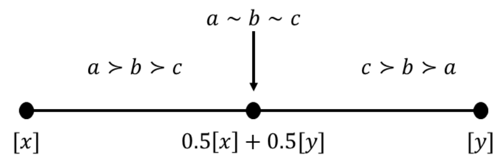

Figure 2.5.1 When relative preference intensities are not unique

u	x	y
a	0	0
b	-1	1
c	-3	3

v	x	y
a	0	0
b	-2	2
c	-3	3

Table 2.5.1 Expected utility representations for the conditional preference relation in Figure 2.5.1

is 2. In terms of preference intensities, it means that the intensity by which the DM prefers b to c at x is twice the intensity by which he prefers a to b at x.

Now turn to the utility function v, which also represents the conditional preference relation. At state x, the expected utility difference between a and b is 2, whereas the expected utility difference between b and c is 1. That is, under the utility function v, the intensity by which the DM prefers a to b at x is twice the intensity by which he prefers b to c at x.

By comparing, we thus see that the relative preference intensities under u are different from those under v. Or, equivalently, the expected utility differences under u differ by more than just a multiplicative constant from those under v.

Another conditional preference relation where the relative preference intensities are not unique is given in Figure 2.5.2.

Question 2.5.1 *Consider the conditional preference relation from Figure 2.5.2.*

(a) Which of the prerequisites in Theorem 2.5.1 is violated?

(b) Give two expected utility representations that induce different relative preference intensities.

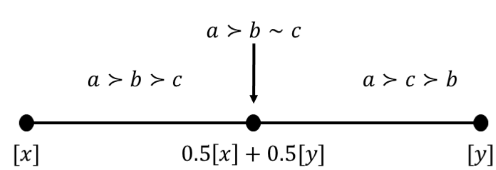

Figure 2.5.2 When relative preference intensities are not unique

	rainy	*stormy*	*calm*
house	10	17	0
garden	0	7	10
tent	3	0	3

Table 2.6.1 Expected utility representation for the conditional preference relation in Figure 2.3.1

These two examples show that the relative preference intensities may no longer be unique if some of the conditions in Theorem 2.5.1 are violated.

2.6 Strict Dominance

Recall from Definition 3.3.1 that a choice is called *rational* if it is optimal for at least one belief. Thus, a choice is *irrational* if it is never optimal for any belief. In this section we will characterize the irrational choices as those that are either strictly dominated by another choice, or strictly dominated by a *randomized* choice. Here, a randomized choice represents a probabilistic mechanism that selects each of your choices with a given probability. A randomized choice is said to *strictly dominate* a choice a if the *expected* intensity by which it is preferred to a is always strictly positive, no matter which belief you hold.

Reconsider the conditional preference relation from Figure 2.3.1. An associated expected utility representation is given by the utility function in Table 2.6.1, which is a copy of Table 2.3.1. As can be seen from Figure 2.3.1, there is no belief for which the choice *tent* is optimal. As such, the choice *tent* is *irrational.* At the same time, your choice *tent* is not strictly dominated by your choice *garden,* as there are beliefs where you prefer *tent* to *garden.* Similarly, your choice *tent* is also not strictly dominated by your choice *house.*

But consider now the situation where you toss a coin to decide whether you will organize the party in the *garden* or in your *house.* This is called a *randomized* choice, that selects your choices *garden* and *house* with probability 0.5. We will see that by using this randomized choice, the *expected* intensity by which you prefer the selected choice to *tent* will always be greater than zero, no matter which belief you hold.

To see this, suppose first that it will be *rainy.* At this state, the utility difference between *house* and *tent* is

$$u(\textit{house, rainy}) - u(\textit{tent, rainy}) = 10 - 3 = 7,$$

whereas the utility difference between *garden* and *tent* is

$$u(\textit{garden, rainy}) - u(\textit{tent, rainy}) = 0 - 3 = -3.$$

Recall that these utility differences may be viewed as the *intensities* by which you prefer the choice *house* or *garden* to *tent.* The intensity by which you prefer *garden* to *tent* at the state *rainy* is negative, which indicates that you prefer *tent* to *garden* when it rains.

As the randomized choice selects both choices *house* and *garden* with probability 0.5, the *expected* intensity by which you prefer the selected choice to *tent* at the state *rainy* is

$$(0.5) \cdot 7 + (0.5) \cdot (-3) = 2 > 0.$$

Similarly, the expected intensity by which you prefer the selected choice to *tent* at the state *stormy* is

$$(0.5) \cdot (17 - 0) + (0.5) \cdot (7 - 0) = 12 > 0,$$

whereas the expected intensity by which you prefer the selected choice to *tent* at the state *calm* is

$$(0.5) \cdot (0 - 3) + (0.5) \cdot (10 - 3) = 2 > 0.$$

Please verify this.

That is, independent of the state of weather, the expected intensity by which you prefer the selected choice to the choice *tent* will alway be greater than zero. This implies, in turn, that for every possible *belief* the expected intensity by which you prefer the selected choice to the choice *tent* will be greater than zero as well. In this case, we say that the randomized choice *strictly dominates* the choice *tent.* In general, strict domination by a randomized choice can be defined as follows.

Definition 2.6.1 (Strict domination by a randomized choice) *A* ***randomized choice*** *is a probability distribution $r \in \Delta(C)$, which assigns to every choice $c \in C$ some probability $r(c)$.*

Consider a conditional preference relation $\succsim$ with an expected utility representation u. A choice a is ***strictly dominated*** *by a randomized choice r under the expected utility representation u if*

$$\sum_{c \in C} r(c) \cdot (u(c,s) - u(a,s)) > 0$$

for every state s.

Here, $u(c,s) - u(a,s)$ reflects the intensity by which the DM prefers the selected choice c to the choice a at the state s. As such, the sum represents the *expected intensity* by which the DM prefers the selected choice to the choice a at the state s if he uses the randomized choice r. If this expected intensity is always greater than zero for every state, then we say that the randomized choice r strictly dominates the choice a.

Let us go back to the conditional preference relation in Figure 2.3.1, with the associated expected utility representation in Table 2.6.1. We have seen that the irrational choice *tent* is strictly dominated by the randomized choice that selects *house* and *garden* with probability 0.5. Moreover, it can be verified that the rational choices *house* and *garden* are not strictly dominated by another choice, nor by a randomized choice. Thus, in this example, the irrational choices are precisely the choices that are strictly dominated by a randomized choice, or by another choice.

The following result shows that this property holds in general: For every conditional preference relation with an expected utility representation, the irrational choices are precisely those that are strictly dominated by a randomized choice, or by another choice. Morover, for every irrational choice a we can always find a randomized choice, or another choice, that strictly dominates *a irrespective of the specific expected utility representation.* In other words, with respect to *every* expected utility representation, the expected intensity by which the DM prefers the selected choice to the choice a will always be greater than zero.

Theorem 2.6.1 (Strict dominance) *Consider a conditional preference relation $\succsim$ with an expected utility representation u.*

(a) A choice a is irrational, if and only if, it is strictly dominated by another choice, or strictly dominated by a randomized choice, under u.

(b) If choice a is strictly dominated by a randomized choice under u, then there is a randomized choice r that strictly dominates choice a under every expected utility representation.

	rainy	*stormy*	*calm*
house	10	17	0
garden	0	7	10
tent	3	0	3
square	-1	-1	-1

Table 2.6.2 When the specific utility representation matters for strict dominance

An immediate consequence of this theorem is that for every choice a we can either find (i) a belief for which a is optimal, or (ii) another choice that strictly dominates a, or (iii) a randomized choice that strictly dominates a. Case (i) applies when a is rational, whereas cases (ii) and (iii) apply when a is irrational. As we will see later in this book, Theorem 2.6.1 will play a very important role when studying games.

To illustrate part (b) in Theorem 2.6.1, let us go back to the conditional preference relation $\succsim$ given by Table 2.6.1. We saw that for this specific expected utility representation u, the irrational choice *tent* is strictly dominated by the randomized choice that selects *house* and *garden* with probability 0.5. But we can say even more: The choice *tent* is strictly dominated by this randomized choice for *every* utility function that represents $\succsim$.

To see this, observe from Figure 2.3.1 that there are preference reversals between every pair of choices, and that there are beliefs where you are indifferent between some, but not all, choices. Therefore, we know by Theorem 2.5.1 that the utility differences are unique up to a positive multiplicative constant. In other words, for every other utility function v that represents $\succsim$, there is a positive number $\alpha > 0$ such that

$$v(a,s) - v(b,s) = \alpha \cdot (u(a,s) - u(b,s))$$

for every state s and every two choices a and b. Here, u is the utility function from Table 2.6.1. Consequently, if *tent* is strictly dominated by the randomized choice above for the utility function u from Table 2.6.1, then this will still be the case for every other expected utility representation v as well. Please verify this.

But there are situations where a choice a is strictly dominated by a given randomized choice r for some expected utility representation u, but not for some other expected utility representation v. As an illustration, consider the conditional preference relation $\succsim$ given by the utility function u in Table 2.6.2. That is, you consider a fourth possible location, *square,* for your birthday party, but you always prefer every other location over *square,* no matter what belief you hold.

Then, your choice *tent* is strictly dominated by the randomized choice r that selects *house, garden* and *square* with probabilities 0.4, 0.4 and 0.2, respectively. Please verify this.

Question 2.6.1 *Find an alternative expected utility representation v such that your choice tent is not strictly dominated by this randomized choice r under v.*

At the same time, there *is* a randomized choice that strictly dominates the irrational choice *tent* for *every* expected utility representation. Consider, for instance, the randomized choice r' that selects *house* and *garden* with probability 0.5. Then, no matter which expected utility representation v we choose, your choice *tent* will always be strictly dominated by r'.

The reason is that the conditional preference relation $\succsim$, when restricted to the choices *house, garden* and *tent,* has preference reversals for every pair of choices, and there are beliefs where you are indifferent between some, but not all, of these three choices Therefore, we can conclude by Theorem

2.5.1 that for every other expected utility representation v, the utility differences for the choices *house, garden* and *tent* will only differ from those under u by a positive multiplicative constant $\alpha > 0$. But then, under every such alternative expected utility representation v, the choice *tent* will still be strictly dominated by the randomized choice r'.

2.7 Proofs

2.7.1 Elements of Linear Algebra and Analysis

In many of the proofs that we will present, we use concepts and results from linear algebra and analysis. In this subsection we give an overview of the elements we use.

2.7.1.1 Linear Spaces

For a number $n \in \mathbf{N}$, we denote by $\mathbf{R}^n$ the set of all n-tuples $v = (v_1, v_2, ..., v_n)$, where $v_1, ..., v_n \in \mathbf{R}$. Such an n-tuple v is called a *vector.* For two vectors $v = (v_1, ..., v_n)$ and $w = (w_1, ..., w_n)$ in $\mathbf{R}^n$, the *sum* of the two vectors is given by $v + w = (v_1 + w_1, ..., v_n + w_n)$, which is again a vector in $\mathbf{R}^n$. For a vector $v = (v_1, ..., v_n)$ and a number $\lambda \in \mathbf{R}$, the *scalar product* $\lambda \cdot v$ is given by $\lambda \cdot v = (\lambda \cdot v_1, ..., \lambda \cdot v_n)$, which is again a vector in $\mathbf{R}^n$.

Now consider a finite set X. By $\mathbf{R}^X$ we denote the set of all functions $v : X \to \mathbf{R}$. Such a function v is again called a *vector.* If $X = \{x_1, ..., x_n\}$, then we can identify the vector $v \in \mathbf{R}^X$ with the vector $(v_1, ..., v_n)$ in $\mathbf{R}^n$, where $v_1 = v(x_1), ..., v_n = v(x_n)$. For instance, if $X = \{x_1, x_2, x_3\}$, and v is the vector in $\mathbf{R}^X$ with $v(x_1) = 2$, $v(x_2) = 5$ and $v(x_3) = -1$, then v can be identified with the vector $(2, 5, -1)$ in $\mathbf{R}^3$.

Note that a probability distribution p on X, as given by Definition 2.2.1, is a special vector in $\mathbf{R}^X$ where $p(x) \geq 0$ for every $x \in X$, and $\sum_{x \in X} p(x) = 1$.

For two vectors v, w in $\mathbf{R}^X$, the *sum* $v + w$ of the two vectors is given by $(v + w)(x) = v(x) + w(x)$ for all $x \in X$, which is again a vector in $\mathbf{R}^X$. For a vector v in $\mathbf{R}^X$ and a number λ in $\mathbf{R}$, the *scalar product* $\lambda \cdot v$ is given by $(\lambda \cdot v)(x) = \lambda \cdot v(x)$ for all $x \in X$, which is again a vector in $\mathbf{R}^X$. The set $\mathbf{R}^X$ together with these two operations, sum and scalar product, constitutes a *linear space.*

By $\underline{0}$ we denote the vector in $\mathbf{R}^X$ where $\underline{0}(x) = 0$ for all $x \in X$. This vector thus corresponds to the zero-vector $(0, ..., 0)$ in $\mathbf{R}^n$.

2.7.1.2 Linear Combinations and Subspaces

Take vectors $v_1, ..., v_m$ in $\mathbf{R}^X$, and numbers $\lambda_1, ..., \lambda_m$ in $\mathbf{R}$. Then, the vector

$$\lambda_1 \cdot v_1 + \lambda_2 \cdot v_2 + ... + \lambda_m \cdot v_m$$

is called a *linear combination* of the vectors $v_1, ..., v_m$.

Now, consider a subset $V \subseteq \mathbf{R}^X$. By $span(V)$ we denote the set of all linear combinations that we can make of vectors in V. More precisely,

$$span(V) := \{\lambda_1 \cdot v_1 + ... + \lambda_m \cdot v_m \mid m \in \mathbf{N},\, v_1, ..., v_m \in V \text{ and } \lambda_1, ..., \lambda_m \in \mathbf{R}\}.$$

We call $span(V)$ the *span* of V. The subset $V \subseteq \mathbf{R}^X$ is called a *linear subspace* of $\mathbf{R}^X$ if $span(V) \subseteq V$. That is, if every linear combination of vectors in V is again in V. In this case, we have in fact that $V = span(V)$.

2.7.1.3 Linear Independence, Basis and Dimension

Consider some vectors $v_1, ..., v_m$ in $\mathbf{R}^X$. We say that $v_1, ..., v_m$ are *linearly independent* if none of the vectors is a linear combination of the other vectors. That is, v_1 cannot be written as

$$v_1 = \lambda_2 \cdot v_2 + ... + \lambda_m \cdot v_m$$

for some $\lambda_2, ..., \lambda_m \in \mathbf{R}$, and the same holds for $v_2, ..., v_m$.

Consider a linear subspace V of $\mathbf{R}^X$, and vectors $v_1, ..., v_m \in V$. The set of vectors $\{v_1, ..., v_m\}$ is a *basis* for V if $v_1, ..., v_m$ are linearly independent, and $span(\{v_1, ..., v_m\}) = V$. That is, every vector in V can be written as a linear combination of the vectors $v_1, ..., v_m$.

Every basis for V has the same number of vectors, and this number is called the *dimension* of V, denoted by $\dim(V)$. If $V = \{\underline{0}\}$, then we set $\dim(V) = 0$.

2.7.1.4 Convex Combinations

Take vectors $v_1, ..., v_m$ in $\mathbf{R}^X$, and numbers $\lambda_1, ..., \lambda_m$ in $\mathbf{R}$ such that $\lambda_1, ..., \lambda_m \geq 0$ and $\lambda_1 + ... + \lambda_m = 1$. Then, the vector

$$\lambda_1 \cdot v_1 + \lambda_2 \cdot v_2 + ... + \lambda_m \cdot v_m$$

is called a *convex combination* of the vectors $v_1, ..., v_m$. For a subset $V \subseteq \mathbf{R}^X$, we denote by $conv(V)$ the set of all convex combinations we can make of vectors in V. That is,

$$conv(V) := \{\lambda_1 \cdot v_1 + ... + \lambda_m \cdot v_m \mid m \in \mathbf{N},\ v_1, ..., v_m \in V \text{ and } \lambda_1, ..., \lambda_m \in \mathbf{R} \\ \text{with } \lambda_1, ..., \lambda_m \geq 0 \text{ and } \lambda_1 + ... + \lambda_m = 1\}.$$

We call $conv(V)$ the *convex hull* of V. The set V is called *convex* if $conv(V) \subseteq V$.

2.7.1.5 Vector Product and Hyperplanes

Consider two vectors v, w in $\mathbf{R}^X$. The *vector product* $v \cdot w$ is the number given by $v \cdot w := \sum_{x \in X} v(x) w(x)$.

A *hyperplane* is a set of the form $H = \{v \in \mathbf{R}^X \mid v \cdot w = c\}$, where w is a vector in $\mathbf{R}^X$ not equal to the zero-vector $\underline{0}$, and c is a number in $\mathbf{R}$. If $c = 0$ then H is a linear subspace of dimension $|X| - 1$, where $|X|$ denotes the number of elements in X. Two hyperplanes H and H' are *parallel* if there is some vector w in $\mathbf{R}^X$ not equal to the zero-vector $\underline{0}$, and two numbers c, c' in $\mathbf{R}$, such that $H = \{v \in \mathbf{R}^X \mid v \cdot w = c\}$ and $H' = \{v \in \mathbf{R}^X \mid v \cdot w = c'\}$.

2.7.2 Proof for Section 2.5

To prove Theorem 2.5.1 we need two preparatory results. In the first result, we denote by $P_{a \sim b}$ the set of beliefs p where $a \sim_p b$.

Lemma 2.7.1 (Span of an indifference set) *Consider a conditional preference relation $\succsim$ that has an expected utility representation, and two choices a and b. Then, $P_{a \sim b}$ is a convex set, and*

$$span(P_{a \sim b}) = \{\lambda_1 \cdot p_1 + \lambda_2 \cdot p_2 \mid p_1, p_2 \in P_{a \sim b} \text{ and } \lambda_1, \lambda_2 \in \mathbf{R}\}.$$

Proof. Suppose that $\succsim$ has expected utility representation u. We first show that $P_{a \sim b}$ is a convex set. To that purpose, take $p_1, ..., p_m \in P_{a \sim b}$ and numbers $\lambda_1, ..., \lambda_m$ with $\lambda_k \geq 0$ for all k and $\sum_{k=1}^m \lambda_k = 1$. We show that $p := \sum_{k=1}^m \lambda_k \cdot p_k \in P_{a \sim b}$. Note that $u(a, p_k) = u(b, p_k)$ for all k, and hence

$$u(a, p) = \sum_{k=1}^m \lambda_k \cdot u(a, p_k) = \sum_{k=1}^m \lambda_k \cdot u(b, p_k) = u(b, p).$$

This means that $p \in P_{a \sim b}$. Thus, $P_{a \sim b}$ is a convex set.

Let

$$A := \{\lambda_1 \cdot p_1 + \lambda_2 \cdot p_2 \mid p_1, p_2 \in P_{a\sim b} \text{ and } \lambda_1, \lambda_2 \in \mathbf{R}\}.$$

We will show that $span(P_{a\sim b}) = A$. Clearly, $A \subseteq span(P_{a\sim b})$. Hence, it remains to show that $span(P_{a\sim b}) \subseteq A$.

Take some $p \in span(P_{a\sim b})$. Then, there are some beliefs $p_1, ..., p_k, p_{k+1}, ..., p_{k+m} \in P_{a\sim b}$ and numbers $\lambda_1, ..., \lambda_k, \lambda_{k+1}, ..., \lambda_{k+m} > 0$ such that

$$p = \lambda_1 p_1 + ... + \lambda_k p_k - \lambda_{k+1} p_{k+1} - ... - \lambda_{k+m} p_{k+m}. \tag{2.7.1}$$

Let $\alpha_1 := \lambda_1 + ... + \lambda_k$ and $\alpha_2 := \lambda_{k+1} + ... + \lambda_{k+m}$. If $\alpha_1 > 0$ and $\alpha_2 > 0$, then define the vectors

$$q_1 := \frac{\lambda_1}{\alpha_1} p_1 + ... + \frac{\lambda_k}{\alpha_1} p_k \text{ and } q_2 := \frac{\lambda_{k+1}}{\alpha_2} p_{k+1} + ... + \frac{\lambda_{k+m}}{\alpha_2} p_{k+m}.$$

It may be verified that q_1 and q_2 are convex combinations of beliefs in $P_{a\sim b}$. Since $P_{a\sim b}$ is a convex set, it follows that $q_1, q_2 \in P_{a\sim b}$. By (2.7.1) we have that

$$p = \alpha_1 q_1 - \alpha_2 q_2,$$

and thus $p \in A$.

If $\alpha_1 > 0$ and $\alpha_2 = 0$, then we must have that $\lambda_{k+1} = ... = \lambda_{k+m} = 0$. We can define $q_1 \in P_{a\sim b}$ as above, and get $p = \alpha_1 q_1$. Thus, $p = \alpha_1 q_1 + 0 \cdot q_1$, which is in A. The case when $\alpha_1 = 0$ and $\alpha_2 > 0$ is similar. Finally, when $\alpha_1 = 0$ and $\alpha_2 = 0$, then $\lambda_1 = ... = \lambda_{k+m} = 0$, which means that $p = \underline{0}$. Thus, $p = 0 \cdot p_1 + 0 \cdot p_2$ for two arbitrary beliefs $p_1, p_2 \in P_{a\sim b}$, and hence $p \in A$.

In general, we thus see that every $p \in span(P_{a\sim b})$ is also in A, and thus $span(P_{a\sim b}) \subseteq A$. Together with the observation above that $A \subseteq span(P_{a\sim b})$, we conclude that $span(P_{a\sim b}) = A$. This completes the proof. ■

The second preparatory result contains some further properties of the set of beliefs where the DM is indifferent between a and b, gathered in Lemma 2.7.2. In this lemma, we denote by $S_{a\sim b}$ the set of states s where $a \sim_{[s]} b$. From now on, we often write $a \sim_s b$ instead of $a \sim_{[s]} b$.

Lemma 2.7.2 (Linear structure of indifference sets) *Suppose there are two choices, a and b, and n states. Consider a conditional preference relation $\succsim$ that has an expected utility representation. Then, the following properties hold:*

(a) $P_{a\sim b} = span(P_{a\sim b}) \cap \Delta(S)$;

(b) if $\succsim$ has preference reversals between a and b, then $span(P_{a\sim b})$ is a hyperplane with dimension $n-1$.

Proof. **(a)** Clearly, $P_{a\sim b} \subseteq span(P_{a\sim b}) \cap \Delta(S)$. It remains to show that $span(P_{a\sim b}) \cap \Delta(S) \subseteq P_{a\sim b}$.

Take some $p \in span(P_{a\sim b}) \cap \Delta(S)$. Then, by Lemma 2.7.1, there are beliefs $p_1, p_2 \in P_{a\sim b}$ and numbers λ_1, λ_2 such that

$$p = \lambda_1 p_1 + \lambda_2 p_2. \tag{2.7.2}$$

Since $p \in \Delta(S)$, we must have that $\sum_{s\in S} p(s) = 1$. Moreover, as p_1, p_2 are beliefs, it holds that $\sum_{s\in S} p_1(s) = \sum_{s\in S} p_2(s) = 1$. In view of (2.7.2),

$$1 = \sum_{s\in S} p(s) = \sum_{s\in S} (\lambda_1 p_1(s) + \lambda_2 p_2(s)) = \lambda_1 \left(\sum_{s\in S} p_1(s)\right) + \lambda_2 \left(\sum_{s\in S} p_2(s)\right) = \lambda_1 + \lambda_2.$$

Suppose first that $\lambda_1 = 0$. Then, $\lambda_2 = 1$, and hence $p = \lambda_2 p_2 = p_2$, which is in $P_{a \sim b}$. The case where $\lambda_2 = 0$ is similar.

Assume next that $\lambda_1, \lambda_2 > 0$. As $\lambda_1 + \lambda_2 = 1$, it follows from (2.7.2) that p is a convex combination of p_1 and p_2, which are both in $P_{a \sim b}$. Since $P_{a \sim b}$ is a convex set by Lemma 2.7.1 , it follows that $p \in P_{a \sim b}$.

Suppose now that $\lambda_1 > 0$ and $\lambda_2 < 0$. Since $\lambda_1 + \lambda_2 = 1$, it must be that $\lambda_1 > 1$. Then, it follows from (2.7.2) that

$$p_1 = \frac{1}{\lambda_1} p - \frac{\lambda_2}{\lambda_1} p_2 = \frac{1}{\lambda_1} p + (1 - \frac{1}{\lambda_1}) p_2 \tag{2.7.3}$$

since $\lambda_2 = 1 - \lambda_1$. As $\lambda_1 > 1$, it follows that p_1 is a convex combination of p and p_2, where p_1 and p_2 are both in $P_{a \sim b}$.

We will show that p must be in $P_{a \sim b}$. Suppose, on the contrary, that $p \notin P_{a \sim b}$. Assume, without loss of generality, that $p \in P_{a \succ b}$, where $P_{a \succ b}$ is the set of beliefs q where $a \succ_q b$. If $\succsim$ is represented by the utility function u, then we have that $u(a, p) > u(b, p)$ and $u(a, p_2) = u(b, p_2)$. By (2.7.3) it then follows that $u(a, p_1) > u(b, p_1)$, which means that $p_1 \in P_{a \succ b}$. This, however, is a contradiction since $p_1 \in P_{a \sim b}$. Hence, we conclude that $p \in P_{a \sim b}$.

The case where $\lambda_1 < 0$ and $\lambda_2 > 0$ is similar. In general, we conclude that every $p \in span(P_{a \sim b}) \cap \Delta(S)$ is also in $P_{a \sim b}$. Hence, $span(P_{a \sim b}) \cap \Delta(S) \subseteq P_{a \sim b}$. As we have already seen that $P_{a \sim b} \subseteq span(P_{a \sim b}) \cap \Delta(S)$, we have that $P_{a \sim b} = span(P_{a \sim b}) \cap \Delta(S)$.

(b) Suppose that $\succsim$ has preference reversals on $\{a, b\}$. Since there is an expected utility representation for $\succsim$, there must be a state x where $a \succ_{[x]} b$, and another state y where $b \succ_{[y]} a$. Moreover, there is a unique belief $p_2 = (1 - \lambda_2)[x] + \lambda_2 [y]$ on the line segment between $[x]$ and $[y]$ where $a \sim_{p_2} b$.

Now, let the remaining states be numbered $s_3, ..., s_n$ such that

$$\begin{gathered} a \succ_{[s_k]} b \text{ for all } k \in \{3, ..., m\}, \\ b \succ_{[s_k]} a \text{ for all } k \in \{m+1, ..., m+l\}, \text{ and} \\ a \sim_{[s_k]} b \text{ for all } k \in \{m+l+1, ..., n\}. \end{gathered}$$

Following the utility design procedure of Section 2.4.1, we choose (i) for every $k \in \{3, ..., m\}$ the unique belief $p_k = (1 - \lambda_k)[s_k] + \lambda_k [y]$ on the line segment between $[s_k]$ and $[y]$ with $a \sim_{p_k} b$, (ii) for every $k \in \{m+1, ..., m+l\}$ the unique belief $p_k = (1 - \lambda_k)[s_k] + \lambda_k [x]$ on the line segment between $[s_k]$ and $[x]$ with $a \sim_{p_k} b$, and (iii) for every $k \in \{m+l+1, ..., n\}$ the belief $p_k = [s_k\}$ with $a \sim_{p_k} b$.

We will now show that $p_2, ..., p_n$ are linearly independent. Take some numbers $\alpha_2, ..., \alpha_n$ such that

$$\sum_{k=2}^{n} \alpha_k \cdot p_k = \underline{0}.$$

By construction, this sum is equal to

$$\begin{aligned}
&\alpha_2((1-\lambda_2)[x]+\lambda_2[y])+\sum_{k=3}^{m}\alpha_k((1-\lambda_k)[s_k]+\lambda_k[y])+\\
&+\sum_{k=m+1}^{m+l}\alpha_k((1-\lambda_k)[s_k]+\lambda_k[x])+\sum_{k=m+l+1}^{n}\alpha_k[s_k]\\
=&\left(\alpha_2(1-\lambda_2)+\sum_{k=m+1}^{m+l}\alpha_k\lambda_k\right)[x]+\left(\alpha_2\lambda_2+\sum_{k=3}^{m}\alpha_k\lambda_k\right)[y]\\
&+\sum_{k=3}^{m+l}\alpha_k(1-\lambda_k)[s_k]+\sum_{k=m+l+1}^{n}\alpha_k[s_k]=\underline{0}.
\end{aligned}$$

As the vectors $[x],[y],[s_3],...,[s_n]$ are linearly independent, and $0<\lambda_k<1$ for all $k\in\{2,...,m+l\}$, it follows that $\alpha_k=0$ for all $k\in\{3,...,n\}$. This, in turn, implies that also $\alpha_2=0$. Hence, the indifference beliefs $p_2,...,p_n\in P_{a\sim b}$ are linearly independent.

As a consequence, the dimension of $span(P_{a\sim b})$ is at least $n-1$. The dimension of $span(P_{a\sim b})$ cannot be n, since otherwise we would have that $span(P_{a\sim b})=\mathbf{R}^S$, and hence, by (a), $P_{a\sim b}=\mathbf{R}^S\cap\Delta(S)=\Delta(S)$. This would contradict the assumption that there are preference reversals between a and b. We thus conclude that the dimension of $span(P_{a\sim b})$ must be $n-1$, and therefore $span(P_{a\sim b})$ is a hyperplane. This completes the proof. ■

Proof of Theorem 2.5.1. Let u,v be two different utility representations for $\succsim$. To prove the statement, we distinguish three cases: (1) there are two choices, (2) there are three choices, and (3) there are at least four choices.

Case 1. Suppose there are two choices, a and b. Since there are preference reversals on $\{a,b\}$, there is some $p^*\in P_{a\succ b}$. Define

$$\alpha:=\frac{v(a,p^*)-v(b,p^*)}{u(a,p^*)-u(b,p^*)}. \tag{2.7.4}$$

We show that

$$v(a,p)-v(b,p)=\alpha\cdot(u(a,p)-u(b,p)) \text{ for all beliefs } p\in\Delta(S). \tag{2.7.5}$$

As there are preference reversals on $\{a,b\}$, it follows by Lemma 2.7.2 (b) that there are $n-1$ linearly independent beliefs $p_1,...,p_{n-1}$ in $P_{a\sim b}$. Moreover, $p^*\notin span(P_{a\sim b})$, as $P_{a\sim b}=span(P_{a\sim b})\cap\Delta(S)$ by Lemma 2.7.2 (a). Hence, $\{p_1,...,p_{n-1},p^*\}$ are linearly independent, and thus form a basis for $\mathbf{R}^S$. As, by construction, $v(a,p_k)-v(b,p_k)=0=\alpha\cdot(u(a,p_k)-u(b,p_k))$ for all $k\in\{1,...,n-1\}$ and, by (2.7.4), $v(a,p^*)-v(b,p^*)=\alpha\cdot(u(a,p^*)-u(b,p^*))$, it follows that (2.7.5) holds for every p in the basis $\{p_1,...,p_{n-1},p^*\}$. Now, take some arbitrary belief $p\in\Delta(S)$. Then, $p=\lambda_1p_1+...+\lambda_{n-1}p_{n-1}+\lambda_np^*$ for some numbers $\lambda_1,...,\lambda_n$. Thus,

$$\begin{aligned}
v(a,p)-v(b,p) &= \sum_{k=1}^{n-1}\lambda_k\cdot(v(a,p_k)-v(b,p_k))+\lambda_n\cdot(v(a,p^*)-v(b,p^*))\\
&= \alpha\cdot\left(\sum_{k=1}^{n-1}\lambda_k\cdot(u(a,p_k)-u(b,p_k))+\lambda_n\cdot(u(a,p^*)-u(b,p^*))\right)\\
&= \alpha\cdot(u(a,p)-u(b,p)),
\end{aligned}$$

which establishes (2.7.5).

Case 2. Suppose there are three choices, a, b and c. Then, by assumption, there is a belief where the DM is indifferent between two, but not all three, choices. Say, there is a belief p where the DM is indifferent between c and b but not between c and a. Then, $P_{c\sim a} \neq P_{c\sim b}$. Let the number α be given by (2.7.4). We show, for every two choices $d, e \in \{a, b, c\}$, that

$$v(d,p) - v(e,p) = \alpha \cdot (u(d,p) - u(e,p)) \text{ for all beliefs } p \in \Delta(S). \tag{2.7.6}$$

By the proof of Case 1, we know that (2.7.6) holds for the choices a and b. We now show that (2.7.6) holds for the choices c and a. Let $p_1, ..., p_{n-1} \in \Delta(S)$ be a basis for $span(P_{c\sim a})$. Then,

$$v(c,p_k) - v(a,p_k) = 0 = \alpha \cdot (u(c,p_k) - u(a,p_k)) \text{ for all } k \in \{1, ..., n-1\}. \tag{2.7.7}$$

Since $P_{c\sim a} \neq P_{c\sim b}$, there is a belief $p_n \in P_{c\sim b} \backslash P_{c\sim a}$. By Lemma 2.7.2 (a) we must then have that $p_n \notin span(P_{c\sim a})$, and hence $\{p_1, ..., p_{n-1}, p_n\}$ is a basis for $\mathbf{R}^S$. As $p_n \in P_{c\sim b}$, it must be that

$$v(c,p_n) - v(b,p_n) = 0 = \alpha \cdot (u(c,p_n) - u(b,p_n)). \tag{2.7.8}$$

Moreover, we know from Case 1 that

$$v(b,p_n) - v(a,p_n) = \alpha \cdot (u(b,p_n) - u(a,p_n)). \tag{2.7.9}$$

If we combine (2.7.8) and (2.7.9), we get

$$\begin{aligned} v(c,p_n) - v(a,p_n) &= (v(c,p_n) - v(b,p_n)) + (v(b,p_n) - v(a,p_n)) \\ &= \alpha \cdot (u(c,p_n) - u(b,p_n)) + \alpha \cdot (u(b,p_n) - u(a,p_n)) \\ &= \alpha \cdot (u(c,p_n) - u(a,p_n)). \end{aligned} \tag{2.7.10}$$

From (2.7.7) and (2.7.10) we conclude, in a similar way as in the proof of Case 1, that

$$v(c,p) - v(a,p) = \alpha \cdot (u(c,p) - u(a,p)) \text{ for all beliefs } p.$$

In a similar fashion we can show (2.7.6) for the choices c and b.

Case 3. Suppose there are at least four choices. By assumption, there are three choices a, b and c with $P_{c\sim a} \neq P_{c\sim b}$. Let the number α be given by (2.7.4). Then, we know by Case 2 that (2.7.6) holds for every $d, e \in \{a, b, c\}$.

We now show (2.7.6) for choices d and a, where d is some arbitrary choice not in $\{a, b, c\}$. To that purpose, we first prove that either $P_{d\sim a} \neq P_{d\sim b}$ or $P_{d\sim a} \neq P_{d\sim c}$. To see this, suppose on the contrary that $P_{d\sim a} = P_{d\sim b} = P_{d\sim c}$. Define $A := P_{d\sim a} = P_{d\sim b} = P_{d\sim c}$. Since $P_{d\sim a} \cap P_{d\sim b} \subseteq P_{a\sim b}$ and $P_{d\sim b} \cap P_{d\sim c} \subseteq P_{b\sim c}$, it follows that $A \subseteq P_{a\sim b}$ and $A \subseteq P_{b\sim c}$. As a consequence, $span(A) \subseteq span(P_{a\sim b})$ and $span(A) \subseteq span(P_{b\sim c})$.

By Lemma 2.7.2 (b) we know that $span(A)$, $span(P_{a\sim b})$ and $span(P_{b\sim c})$ all have dimension $n-1$, and hence it must be that $span(A) = span(P_{a\sim b})$ and $span(A) = span(P_{b\sim c})$. By Lemma 2.7.2 (a) it then follows that $A = P_{a\sim b} = P_{b\sim c}$.

Note that $A = P_{a\sim b} \cap P_{b\sim c} \subseteq P_{a\sim c}$. In a similar way as above, it can be shown that, in fact, $A = P_{a\sim c}$. We thus conclude that $P_{a\sim b} = P_{b\sim c} = P_{a\sim c}$. This is a contradiction to our assumption that $P_{a\sim c} \neq P_{b\sim c}$. Hence, either $P_{d\sim a} \neq P_{d\sim b}$ or $P_{d\sim a} \neq P_{d\sim c}$.

Assume, without loss of generality, that $P_{d\sim a} \neq P_{d\sim b}$. Then, it can be shown in a similar way as for Case 2 that (2.7.6) holds for the choices d and a.

Now, take some choice $d \notin \{a, b, c\}$, and some arbitrary choice $e \notin \{a, d\}$. Since we know that (2.7.6) holds for the choices d and a, and for the choices e and a, it follows that

$$v(d, p) - v(a, p) = \alpha \cdot (u(d, p) - u(a, p)) \text{ for all beliefs } p$$

and

$$v(a, p) - v(e, p) = \alpha \cdot (u(a, p) - u(e, p)) \text{ for all beliefs } p.$$

This implies that

$$\begin{aligned} v(d, p) - v(e, p) &= (v(d, p) - v(a, p)) + (v(a, p) - v(e, p)) \\ &= \alpha \cdot (u(d, p) - u(a, p)) + \alpha \cdot (u(a, p) - u(e, p)) \\ & \quad \alpha \cdot (u(d, p) - u(e, p)) \text{ for all beliefs } p. \end{aligned}$$

Hence, (2.7.6) holds for every two choices d, e. This completes the proof. ■

2.7.3 Proof for Section 2.6

To prove Theorem 2.6.1 we rely on a well-known result in mathematics, called the *separating hyperplane theorem.* To formally state this result we need some additional definitions. For every two vectors $v, w \in \mathbf{R}^X$ we denote by

$$d(v, w) := \sqrt{\sum_{x \in X} (v(x) - w(x))^2}$$

the *Euclidean distance* between v and w. For some vector $v \in \mathbf{R}^X$ and number $r > 0$, we denote by

$$B_r(v) := \{w \in \mathbf{R}^X \mid d(v, w) < r\}$$

the *open ball with radius* r around the vector v. A subset $A \subseteq \mathbf{R}^X$ is called *open* if for every vector $v \in A$ there is some number $r > 0$ such that $B_r(v) \subseteq A$. A subset $A \subseteq \mathbf{R}^X$ is called *closed* if its complement $\mathbf{R}^X \backslash A$ is open. Intuitively, this means that the boundary of A is included in A itself. The set A is called *bounded* if there is some vector $v \in \mathbf{R}^X$ and some $r > 0$ such that $A \subseteq B_r(v)$. Finally, two sets A, B are *disjoint* if the set $A \cap B$ is empty.

The separating hyperplane theorem states that under certain conditions, two disjoints sets in $\mathbf{R}^X$ can be separated by a hyperplane.

Theorem 2.7.1 (Separating hyperplane theorem) *Let X be a finite set, and let A and B be two disjoint, closed and convex sets in $\mathbf{R}^X$, where A is bounded. Then, there is a vector $w \in \mathbf{R}^X \backslash \{\underline{0}\}$ and a number α such that*

$$a \cdot w > \alpha \textit{ for all } a \in A \textit{ and } b \cdot w < \alpha \textit{ for all } b \in B.$$

Recall that the set $H := \{v \in \mathbf{R}^X \mid v \cdot w = \alpha\}$ is a hyperplane. Hence, under the conditions above, the sets A and B are on different sides of the hyperplane H, and are thus *separated* by the hyperplane H.

We are now ready to prove Theorem 2.6.1.

Proof of Theorem 2.6.1. **(a)** Suppose first that a choice a is strictly dominated by another choice. Then, clearly, a is irrational. Assume next that a is strictly dominated by a randomized choice r under the expected utility representation u. We show that a must be irrational.

To see this, note that for every state s we have that

$$\sum_{c \in C} r(c) \cdot (u(c,s) - u(a,s)) > 0.$$

This implies that for every belief p,

$$\sum_{c \in C} r(c) \cdot (u(c,p) - u(a,p)) > 0.$$

Consequently, for every belief p there is some choice c with $r(c) > 0$ such that $u(c,p) > u(a,p)$. In particular, the choice a cannot be optimal for any belief, and thus a is irrational.

Next, suppose that choice a is irrational. We will show that a is either strictly dominated by another choice, or strictly dominated by a randomized choice under the expected utility representation u.

Let the set of choices be $C = \{c_1, ..., c_m\}$. For every belief p, consider the utility difference vector $d[p]$ in $\mathbf{R}^C$ where

$$d[p](c) := u(c,p) - u(a,p) \text{ for every } c \in C.$$

Thus, $d[p]$ specifies for every choice c the intensity by which the DM prefers the choice c to the fixed choice a. Let A be the subset of $\mathbf{R}^C$ given by

$$A := \{d[p] \mid p \in \Delta(S)\}.$$

That is, A contains all utility difference vectors with respect to choice a.

Moreover, define the set

$$B = \{v \in \mathbf{R}^C \mid v(c) \leq 0 \text{ for all } c \in C\}.$$

We will now show that the sets A and B satisfy the conditions in Theorem 2.7.1.

We first show that A and B are disjoint. To see this, take an arbitrary vector $d[p]$ in A. As a is irrational, we know that a is not optimal for the belief p, and hence there is some choice c with $d[p](c) = u(c,p) - u(a,p) > 0$. Thus, $d(p)$ is not in B. Therefore, A and B are disjoint.

It may easily be verified that A is closed and bounded, and that B is closed and convex. It remains to show that A is convex. Take some vectors $d[p], d[q]$ in A and some number $\lambda \in [0,1]$. We will show that $(1-\lambda) \cdot d[p] + \lambda \cdot d[q]$ is in A. For every choice c we have that

$$\begin{aligned}((1-\lambda) \cdot d[p] + \lambda \cdot d[q])(c) &= (1-\lambda) \cdot (u(c,p) - u(a,p)) + \lambda \cdot (u(c,q) - u(a,q)) \\ &= u(c, (1-\lambda)p + \lambda q) - u(a, (1-\lambda)p + \lambda q) = d[(1-\lambda)p + \lambda q](c).\end{aligned}$$

As $(1-\lambda)p + \lambda q$ is again a belief, we conclude that $(1-\lambda) \cdot d[p] + \lambda \cdot d[q]$ is in A. Thus, A is convex.

By Theorem 2.7.1 we thus conclude that there is some vector $w \in \mathbf{R}^C \backslash \{\underline{0}\}$ and a number α such that

$$a \cdot w > \alpha \text{ for all } a \in A \text{ and } b \cdot w < \alpha \text{ for all } b \in B. \tag{2.7.11}$$

As $\underline{0} \in B$, it follows that $\alpha > \underline{0} \cdot w = 0$, and hence $a \cdot w > 0$ for all $a \in A$.

We next show that $w(c) \geq 0$ for every $c \in C$. Suppose, contrary to what we want to show, that $w(c^*) < 0$ for some choice c^*. Then, the vector $b \in \mathbf{R}^C$ with $b(c^*) = \frac{\alpha}{w(c^*)} < 0$ and $b(c) = 0$ for all $c \neq c^*$ is in B. Moreover, $b \cdot w = b(c^*) \cdot w(c^*) = \alpha$, which is a contradiction since $b \cdot w < \alpha$ for all $b \in B$. Thus, $w(c) \geq 0$ for all $c \in C$.

We now show that $w(c) > 0$ for some choice c. Suppose, on the contrary, that $w(c) = 0$ for all $c \in C$. Then, $a \cdot w = 0 < \alpha$ for all $a \in A$, which would be a contradiction to (2.7.11).

Summarizing, we see that $w(c) \geq 0$ for all choices c, and $w(c) > 0$ for some choice c. Now, define the randomized choice r by

$$r(c) := \frac{w(c)}{\sum_{d \in C} w(d)} \text{ for all } c \in C.$$

As $w(c) \geq 0$ for all choices c, and $w(c) > 0$ for some choice c, it follows that $r(c) \geq 0$ for every $c \in C$, and $\sum_{c \in C} r(c) = 1$. That is, r is a well-defined randomized choice.

For every belief p we then have that

$$\begin{aligned} \sum_{c \in C} r(c) \cdot (u(c,p) - u(a,p)) &= \frac{1}{\sum_{d \in C} w(d)} \cdot \sum_{c \in C} w(c) \cdot (u(c,p) - u(a,p)) \\ &= \frac{1}{\sum_{d \in C} w(d)} \cdot (w \cdot d[p]) > 0. \end{aligned}$$

Here, the inequality follows from the insight in (2.7.11) that $w \cdot a > 0$ for all $a \in A$, and the fact that $d[p] \in A$. Hence, we conclude that

$$\sum_{c \in C} r(c) \cdot (u(c,p) - u(a,p)) > 0 \text{ for all beliefs } p,$$

which means that the choice a is strictly dominated by the randomized choice r under u. This completes the proof for part (a).

(b) Suppose now that a is strictly dominated by a randomized choice r' under u. We will show that there is a randomized choice r that strictly dominates a under *every* utility function v that represents $\succsim$.

We start by assuming that there are no equivalent choices. At the end of the proof, we show how to prove the statement if there are equivalent choices.

Let D be the set of choices that are not weakly dominated by any other choice. We first show that D contains at least one choice. Suppose not. Then, every choice c in C would be weakly dominated by some other choice $f(c)$ in C. Start from some arbitrary choice c. Then, c is weakly dominated by $f(c)$, $f(c)$ is weakly dominated by $f(f(c))$, and so on, without end. This induces an infinite sequence of choices

$$c^1 \to c^2 \to c^3 \to ...$$

where $c^1 = c$, $c^2 = f(c^1)$, $c^3 = f(c^2)$, and so on. As there are only finitely many choices, there must be a choice that occurs at least twice in this sequence. That is, there must be some $k, m \geq 1$ such that

$$c^k \to c^{k+1} \to ... \to c^{k+m} = c^k. \tag{2.7.12}$$

Hence, c^k weakly dominates c^{k+m-1}, which in turn weakly dominates c^{k+m-2}, and so on, until c^{k+1} which weakly dominates c^k. However, if a choice b weakly dominates c, and c weakly dominates d, then it follows that b weakly dominates d. Please verify this. If we apply this insight repeatedly to (2.7.12), then we conclude that c^k weakly dominates c^k. This, however, cannot be. Please verify this. Thus, the assumption above that D contains no choice cannot be true. That is, D must contain at least one choice.

We next show that every choice not in D must be weakly dominated by some choice in D. To see this, consider a choice $c^1 \notin D$. Then, there is a choice c^2 that weakly dominates c^1. If $c^2 \in D$ we are done. If $c^2 \notin D$, then there is some choice c^3 that weakly dominates c^2, and so on. We thus obtain a sequence of choices

$$c^1 \to c^2 \to ...$$

This sequence, however, cannot be infinite, since in that case we would conclude that some choice is weakly dominated by itself, like above. Thus, there must be some c^k where the sequence terminates, and hence we obtain

$$c^1 \to c^2 \to ... \to c^k.$$

As c^k is not weakly dominated by any other choice, c^k must be in D. By the same argument as above, it can then be concluded that c^k weakly dominates c^1. Thus, every choice $c^1 \notin D$ is weakly dominated by some choice $c^k \in D$.

We now show that for every belief p there is a choice in D that is optimal for p. To see this, consider a belief p and a choice c that is optimal for p. If c is in D, then we are done. If c is not in D, then we know from above that there is some $d \in D$ that weakly dominates c. Since $u(d,p) \geq u(c,p)$ and c is optimal for the belief p, it follows that d is optimal for the belief p as well.

Recall our assumption above that a is strictly dominated by some randomized choice under u. Hence, by part (a) of Theorem 2.6.1 we know that a is irrational under $\succsim$. That is, a is not optimal for any belief p. Since we have seen above that for every belief p there is a choice in D that is optimal for p, it follows that there is no belief p such that a is optimal among all choices in $D \cup \{a\}$. Therefore, a is irrational if we restrict to choices in $D \cup \{a\}$ only.

We will now show that there is a randomized choice r that only assigns positive probability to choices in $D \backslash \{a\}$, and that strictly dominates a under *every* utility function v that represents $\succsim$.

We distinguish two cases: (1) a is in D, and (2) a is not in D.

Case 1. Suppose that a is in D. Then, no choice in $D \cup \{a\}$ weakly dominates another choice in $D \cup \{a\}$. Since we assume that there are no equivalent choices, it follows that there are preference reversals between every two choices in $D \cup \{a\}$.

Note that $D \cup \{a\}$ must contain at least three choices. To see this, assume that $D \cup \{a\}$ would contain only one choice, which would be a. This would be a contradiction, as a is irrational in $D \cup \{a\}$. If $D \cup \{a\}$ would contain two choices, say a and b, then b would strictly dominate a, since a is irrational in $D \cup \{a\}$. This would also be a contradiction, as a is in D, and hence a is not weakly dominated by any other choice. Therefore, $D \cup \{a\}$ must contain at least three choices.

We next show that there are choices $b, c, d \in D$ with $P_{b\sim d} \neq P_{c\sim d}$. Suppose not. Then, in particular, $P_{a\sim b} = P_{a\sim c}$ for every $b, c \in D \backslash \{a\}$. Take some choice $b \in D \backslash \{a\}$ and some belief $p \in P_{a\sim b}$. Then, $p \in P_{a\sim c}$ for every $c \in D \backslash \{a\}$. In particular, a would be optimal for the belief p among the choices in $D \cup \{a\}$. This would be a contradiction, as a is irrational among the choices in $D \cup \{a\}$. Thus, there are choices $b, c, d \in D$ with $P_{b\sim d} \neq P_{c\sim d}$. Hence, there are beliefs where the DM is indifferent between some, but not all, choices in $D \cup \{a\}$.

Since we also know that there are preference reversals between any two choices in $D \cup \{a\}$, it follows from Theorem 2.5.1 that for every other expected utility representation v there is some number $\alpha > 0$ with

$$v(c,s) - v(a,s) = \alpha \cdot (u(c,s) - u(a,s)) \tag{2.7.13}$$

for all choices $c \in D \backslash \{a\}$ and all states s.

Recall that choice a is irrational among choices in $D \cup \{a\}$. By part (a) of this theorem we know that there is a randomized choice r that only assigns positive probability to choices in $D \backslash \{a\}$, and that strictly dominates a under u. As such,

$$\sum_{c \in C} r(c) \cdot (u(c,s) - u(a,s)) > 0 \text{ for all states } s. \tag{2.7.14}$$

By combining (2.7.13) and (2.7.14), and using the fact that the randomized choice r only assigns positive probability to choices in $D\backslash\{a\}$, it follows for every state s that

$$\sum_{c \in C} r(c) \cdot (v(c,s) - v(a,s)) = \alpha \cdot \sum_{c \in C} r(c) \cdot (u(c,s) - u(a,s)) > 0.$$

Thus, r strictly dominates a for every utility function v that represents $\succsim$. This completes the proof of Case 1.

Case 2. Suppose that a is not in D. Let E be the set of choices in D that weakly dominate a, and let F be the set of choices in D that do not weakly dominate a. Since a is not in D, we know that a is weakly dominated by some other choice. Hence, from one of our arguments above it follows that there must be some $d \in D$ that weakly dominates a, and thus E is non-empty. We distinguish two subcases: (2.1) for every state s there is some $e \in E$ with $e \succ_{[s]} a$, and (2.2) there is some state s such that $e \sim_{[s]} a$ for every choice $e \in E$.

Case 2.1. Suppose that for every state s there is some choice $e(s) \in E$ with $e \succ_{[s]} a$. Let r be a randomized choice that assigns positive probability to all choices in E, and probability zero to all other choices. We show that r strictly dominates a under every utility function v that represents $\succsim$.

Let v be some arbitrary utility function that represents $\succsim$. Then, $v(e(s),s) > v(a,s)$ for every state s, and $v(e,s) \geq v(a,s)$ for every state s and every choice $e \in E$. The last inequality holds because a is weakly dominated by every choice in E. But then, for every state s,

$$\sum_{e \in E} r(e) \cdot (v(e,s) - v(a,s)) > 0$$

because r assigns positive probability to all choices in E. Thus, r strictly dominates a under every utility function v that represents $\succsim$. This completes the proof of Case 2.1.

Case 2.2. Suppose there is some state s^* with $e \sim_{[s^*]} a$ for every choice $e \in E$. As a is irrational among choices in $D \cup \{a\}$, we know there must be some $f \in D$ with $f \succ_{[s^*]} a$. Hence, necessarily, f must be in F. In particular, F is non-empty.

Since a is irrational among choices in $D \cup \{a\}$, we know from part (a) of this theorem that there is a randomized choice r that strictly dominates a under u, and that only assigns positive probability to choices in D. We show that r strictly dominates a under every utility function v that represents $\succsim$.

We distinguish two subcases: (2.2.1) either D contains at most two choices or there are choices $b,c,d \in D$ with $P_{b\sim c} \neq P_{b\sim d}$, and (2.2.2) D contains more than two choices and $P_{b\sim c} = P_{b\sim d}$ for all choices $b,c,d \in D$.

Case 2.2.1. Suppose that either D contains at most two choices, or there are beliefs where the DM is indifferent between some, but not all, choices in D. By construction, no choice in D weakly dominates another choice in D, and hence there are preference reversals between all choices in D. As such, we can conclude from Theorem 2.5.1 that for every other expected utility representation v there is some $\alpha > 0$ with

$$v(b,s) - v(c,s) = \alpha \cdot (u(b,s) - u(c,s)) \text{ for all choices } b,c \in D \tag{2.7.15}$$

and all states s.

Take some arbitrary utility function v that represents $\succsim$. Consider the state s^* above, and the choice $f \in F$ with $f \succ_{[s^*]} a$. As, by definition of F, we have that f does not weakly dominate a, and since a does not weakly dominate f, we know that there are preference reversals between a and f. By Lemma 2.7.2 (b), $span(P_{a\sim f})$ has dimension $n-1$, where n is the number of states. Consider a basis $\{p_1, ..., p_{n-1}\}$ for $span(P_{a\sim f})$, consisting of beliefs. Moreover, since $[s^*] \notin P_{a\sim f}$ we know from Lemma 2.7.2 (a) that $[s^*] \notin span(P_{a\sim f})$, and therefore $\{p_1, ..., p_{n-1}, [s^*]\}$ is a basis for $\mathbf{R}^S$.

Since $p_1, ..., p_{n-1} \in P_{a\sim f}$ we must have that

$$v(a, p_k) = v(f, p_k) \text{ for all } k \in \{1, ..., n-1\}.$$

Take some $e \in E$. As $[s^*] \in P_{a\sim e}$, it must be that

$$v(a, s^*) = v(e, s^*).$$

Since $u(a, p_k) = u(f, p_k)$ for all $k \in \{1, ..., n-1\}$, and $u(a, s^*) = u(e, s^*)$, we know in particular that

$$v(a, p_k) - v(f, p_k) = 0 = \alpha \cdot (u(a, p_k) - u(f, p_k)) \text{ for all } k \in \{1, ..., n-1\} \tag{2.7.16}$$

and

$$v(a, s^*) - v(e, s^*) = 0 = \alpha \cdot (u(a, s^*) - u(e, s^*)). \tag{2.7.17}$$

Now, take some arbitrary choice $c \in D$ and some arbitrary belief p. Recall that $\{p_1, ..., p_{n-1}, [s^*]\}$ is a basis for $\mathbf{R}^S$, and hence there are numbers $\lambda_1, ..., \lambda_n$ with $p = \lambda_1 p_1 + ... + \lambda_{n-1} p_{n-1} + \lambda_n [s^*]$. In view of (2.7.15), (2.7.16) and (2.7.17) we then conclude that

$$\begin{aligned} v(a, p) - v(c, p) &= \sum_{k=1}^{n-1} \lambda_k \cdot (v(a, p_k) - v(c, p_k)) + \lambda_n \cdot (v(a, s^*) - v(c, s^*)) \\ &= \sum_{k=1}^{n-1} \lambda_k \cdot (v(a, p_k) - v(f, p_k) + v(f, p_k) - v(c, p_k)) + \\ &\quad + \lambda_n \cdot (v(a, s^*) - v(e, s^*) + v(e, s^*) - v(c, s^*)) \\ &= \alpha \cdot \sum_{k=1}^{n-1} \lambda_k \cdot (u(a, p_k) - u(f, p_k) + u(f, p_k) - u(c, p_k)) + \\ &\quad + \alpha \cdot \lambda_n \cdot (u(a, s^*) - u(e, s^*) + u(e, s^*) - u(c, s^*)) \\ &= \alpha \cdot (u(a, p) - u(c, p)). \end{aligned} \tag{2.7.18}$$

Recall that the randomized choice r strictly dominates a under u. Hence, for every state s we have that

$$\sum_{c \in C} r(c) \cdot (u(c, s) - u(a, s)) > 0.$$

Together with (2.7.18) we conclude, for every state s, that

$$\sum_{c \in C} r(c) \cdot (v(c, s) - v(a, s)) = \alpha \cdot \sum_{c \in C} r(c) \cdot (u(c, s) - u(a, s)) > 0.$$

Thus, r strictly dominates a under v. This completes the proof of Case 2.2.1.

Case 2.2.2. Suppose that D contains more than two choices and there are no beliefs where the DM is indifferent between some, but not all, choices in D. Then, $P_{b\sim c} = P_{d\sim e}$ for all pairs of choices $\{b, c\}$ and $\{d, e\}$ in D. Define $P := P_{b\sim c}$ for some $b, c \in D$.

Now, consider some choice $f \in F$. As f does not weakly dominate a and a does not weakly dominate f, we know that $P_{a\sim f}$ is non-empty. We show that no $p \in P_{a\sim f}$ can be in P. Suppose, on the contrary, that $p \in P_{a\sim f}$ would also be in P. Then, for every choice $c \in D\backslash\{f\}$, we would have that $p \in P_{a\sim f} \cap P_{f\sim c}$. Hence, by transitivity, it would follow that $p \in P_{a\sim c}$ for every $c \in D$. But then, a would be optimal for the belief p among choices in $D \cup \{a\}$, which is a contradiction, since we have seen that a is irrational among choices in $D \cup \{a\}$. We thus conclude that no $p \in P_{a\sim f}$ is in P. In particular, $P_{a\sim f} \neq P$.

By construction of the set F, no choice in $F \cup \{a\}$ weakly dominates another choice in $F \cup \{a\}$. Thus, there are preference reversals between every two choices in $F \cup \{a\}$. Moreover, since $P_{a\sim f} \neq P$, we know that either F only has the choice f, or for every choice $c \in F\backslash\{f\}$ we have that $P_{a\sim f} \neq P_{c\sim f}$. Thus, we conclude by Theorem 2.5.1 that for every other expected utility representation v there is some $\alpha > 0$ with

$$v(b, s) - v(c, s) = \alpha \cdot (u(b, s) - u(c, s)) \text{ for all choices } b, c \in F \cup \{a\} \tag{2.7.19}$$

and all states s.

Take some arbitrary utility function v that represents $\succsim$. We show that the randomized choice r strictly dominates a under v.

Consider some arbitrary choice $e \in E$. Since $e, f \in D$, we know that e does not weakly dominate f, and f does not weakly dominate e, and therefore there are preference reversals between e and f. By Lemma 2.7.2 (b), we then know that $span(P) = span(P_{e\sim f})$ has dimension $n-1$. Let $\{p_1, ..., p_{n-1}\}$ be a basis for $span(P)$ consisting of beliefs in P.

Consider now some $p_n \in P_{a\sim e}$. In the same way as above, it can be shown that p_n is not in P. Hence, by Lemma 2.7.2 (a) we know that $p_n \notin span(P)$. As such, $\{p_1, ..., p_{n-1}, p_n\}$ is a basis for $\mathbf{R}^S$.

Since $p_1, ..., p_{n-1}$ are in $P = P_{e\sim f}$, we must have that

$$v(e, p_k) = v(f, p_k) \text{ for all } k \in \{1, ..., n-1\}.$$

Moreover, as $p_n \in P_{a\sim e}$, it must be that

$$v(e, p_n) = v(a, p_n).$$

As $u(e, p_k) = u(f, p_k)$ for all $p_k \in \{1, .., , n-1\}$ and $u(e, p_n) = u(a, p_n)$, we know that

$$v(e, p_k) - v(f, p_k) = 0 = \alpha \cdot (u(e, p_k) - u(f, p_k)) \text{ for all } k \in \{1, .., , n-1\}, \tag{2.7.20}$$

and

$$v(e, p_n) - v(a, p_n) = 0 = \alpha \cdot (u(e, p_n) - u(a, p_n)). \tag{2.7.21}$$

Take some $e \in E$ and some belief p. Since $\{p_1, ..., p_n\}$ is a basis for $\mathbf{R}^S$, there are numbers $\lambda_1, ..., \lambda_n$

such that $p = \lambda_1 p_1 + ... + \lambda_n p_n$. In view of (2.7.19), (2.7.20) and (2.7.21), we conclude that

$$\begin{aligned} v(e,p) - v(a,p) &= \sum_{k=1}^{n-1} \lambda_k \cdot (v(e,p_k) - v(a,p_k)) + \lambda_n \cdot (v(e,p_n) - v(a,p_n)) \\ &= \sum_{k=1}^{n-1} \lambda_k \cdot (v(e,p_k) - v(f,p_k) + v(f,p_k) - v(a,p_k)) \\ &\quad + \lambda_n \cdot (v(e,p_n) - v(a,p_n)) \\ &= \alpha \cdot \sum_{k=1}^{n-1} \lambda_k \cdot (u(e,p_k) - u(f,p_k) + u(f,p_k) - u(a,p_k)) \\ &\quad + \alpha \cdot \lambda_n \cdot (u(e,p_n) - u(a,p_n)) \\ &= \alpha \cdot (u(e,p) - u(a,p)). \end{aligned}$$

As this holds for every $e \in E$, it follows together with (2.7.19) that

$$v(c,p) - v(a,p) = \alpha \cdot (u(c,p) - u(a,p)) \text{ for all } c \in D. \tag{2.7.22}$$

Since the randomized choice r strictly dominates a under u we know, for every state s, that

$$\sum_{c \in D} r(c) \cdot (u(c,s) - u(a,s)) > 0.$$

Together with (2.7.22) it follows that

$$\sum_{c \in D} r(c) \cdot (v(c,s) - v(a,s)) = \alpha \cdot \sum_{c \in D} r(c) \cdot (u(c,s) - u(a,s)) > 0.$$

Thus, the randomized choice r strictly dominates a under v. This completes the proof of Case 2.2.2.

Since we have exhausted all the cases, we have shown, in the absence of equivalent choices, that there is a randomized choice r that strictly dominates a under every utility function v that represents $\succsim$.

Suppose now that some choices would be equivalent to other choices. We can first reduce C to a subset C^* of choices that does not contain equivalent choices by eliminating, for every set of equivalent choices, all choices but one. Let $\succsim^*$ be the restriction of the conditional preference relation $\succsim$ to choices in C^*. We can do the eliminations in such a way that the irrational choice a we started from is still in C^*.

Clearly, a will still be irrational for $\succsim^*$. By the proof above, we know that there is a randomized choice r on C^* that strictly dominates a for every utility function v^* that represents $\succsim^*$. But then, the same randomized choice r will then strictly dominate a for every utility function v that represents $\succsim$. This completes the proof. ■

Solutions to In-Chapter Questions

Question 2.1.1. The choices *garden* and *tent* are optimal. To see this, note that *garden* is weakly preferred to the other two choices, and *tent* is weakly preferred to the other two choices.

Question 2.1.2. (a) Suppose that $a \succ b$ and $b \succ c$. By transitivity, it follows that $a \succsim c$. Suppose, contrary to what we want to show, that $a \sim c$. Then, $c \sim a$ and $a \succ b$, and hence by transitivity it would follow that $c \succsim b$, which is a contradiction. Thus, $a \succ c$.

(b) Suppose that $a \succ b$ and $b \sim c$. By transitivity, it follows that $a \succsim c$. Suppose, contrary to what we want to show, that $a \sim c$. Then, as $b \sim c$ and $c \sim a$, it would follow by transitivity that $b \succsim a$, which is a contradiction. Thus, $a \succ c$.

(c) Suppose that $a \sim b$ and $b \succ c$. By transitivity, it follows that $a \succsim c$. Suppose, contrary to what we want to show, that $a \sim c$. Then, as $c \sim a$ and $a \sim b$, it would follow by transitivity that $c \succsim b$, which is a contradiction. Thus, $a \succ c$.

(d) Suppose that $a \sim b$ and $b \sim c$. Then, by transitivity, it follows that $a \succsim c$. Moreover, as $c \sim b$ and $b \sim a$, it follows by transitivity that also $c \succsim a$. But then, it must be that $a \sim c$.

Question 2.1.3. (a) The choice b is optimal, because it is weakly preferred to all other choices.

(b) Note that $c \succ d$ and $d \succ a$ but $c \sim a$. This violates property (a) in Question 2.1.2, and thus $\succsim$ is not transitive.

Question 2.1.4. The preference relation $\succsim$ induced by the utility function u is given by $b \succ a$, $c \succ a$, $d \succ a$, $c \succ b$, $b \sim d$ and $c \succ d$.

Question 2.1.5. Take choices a, b, c where $a \succsim b$ and $b \succsim c$. Then, $u(a) \geq u(b)$ and $u(b) \geq u(c)$, which implies that $u(a) \geq u(c)$. Hence, $a \succsim c$. Therefore, $\succsim$ is transitive.

Question 2.1.6. (a) Take $a, b \in C^1$. As both a and b are optimal, we must have that $a \succsim b$ and $b \succsim a$, which implies that $a \sim b$.

(b) Take $a \in C^1$ and $b \notin C^1$. As a is optimal we must have that $a \succsim b$. Assume, contrary to what we want to show, that $a \sim b$. As a is optimal, we know that a is weakly preferred to every other choice. Since $a \sim b$, it would follow from Question 2.1.2 that also b is weakly preferred to all other choices, and hence also b would be optimal. This is a contradiction. Hence, we conclude that $a \succ b$.

Question 2.1.7. It may be verified that for all choices e, f, g with $e \succsim f$ and $f \succsim g$ it holds that $e \succsim g$. Hence, $\succsim$ is transitive.

To find a utility representation we follow the procedure in the proof of Theorem 2.1.2. The set of optimal choices in C is $C^1 = \{b, c\}$. Thus, $C \backslash C^1 = \{a, d\}$. Within $C \backslash C^1$, the set of optimal choices is $C^2 = \{a\}$. Hence, $C^3 = \{d\}$. Choose $u^1 > u^2 > u^3$, and set $u(b) = u(c) := u^1$, $u(a) := u^2$ and $u(d) := u^3$.

Question 2.2.1. (a) The belief q assigns probability 0.4 to *rainy,* probability 0 to *stormy,* and probability 0.6 to *calm.* Thus, $q = (0.4, 0, 0.6)$.

(b) The location of belief $r = (0.6, 0.3, 0.1)$ can be found in Figure 2.7.1.

Question 2.3.1. The expected utility of the three choices are given by

$$\begin{aligned} u(\textit{house, } q) &= (0.3) \cdot 10 + (0.1) \cdot 17 + (0.6) \cdot 0 = 4.7, \\ u(\textit{garden, } q) &= (0.3) \cdot 0 + (0.1) \cdot 7 + (0.6) \cdot 10 = 6.7, \\ u(\textit{tent, } q) &= (0.3) \cdot 3 + (0.1) \cdot 0 + (0.6) \cdot 3 = 2.7. \end{aligned}$$

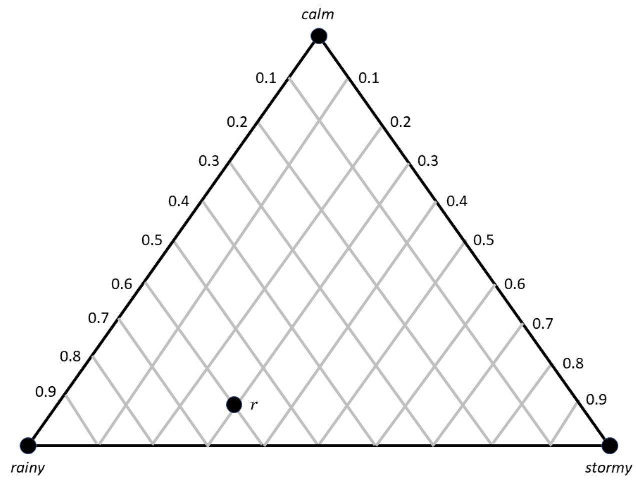

Figure 2.7.1 Question 2.2.1

This matches your preferences at the belief q, which are given by *garden* $\succ_q$ *house* $\succ_q$ *tent.*

Question 2.3.2. Take the expected utility representation from Table 2.3.1 as a starting point. Currently, the utility difference between *garden* and *house* at the state *calm* is 10, but it should become 1. To make this possible, we multiply all utilities in Table 2.3.1 by 1/10, and get

	rainy	*stormy*	*calm*
house	1	1.7	0
garden	0	0.7	1
tent	0.3	0	0.3

To make the three utilities in the first row 0, we substract 1 from all entries in the first column, and substract 1.7 from all entries in the second column, to get

	rainy	*stormy*	*calm*
house	0	0	0
garden	-1	-1	1
tent	-0.7	-1.7	0.3

Question 2.4.1. The conditional preference relation between *garden* and *tent* is depicted in Figure 2.7.2. As requested, we choose the utilities for *tent* at the three states equal to 0, that is, $u(tent,\ rainy) = 0$, $u(tent,\ stormy) = 0$ and $u(tent,\ calm) = 0$. Also, we choose $u(garden,\ calm) = 1$. Consider the belief $p_1 = (0.7, 0, 0.3)$ in Figure 2.7.2 on the line segment between *rainy* and *calm.* By the utility difference property,

$$\frac{u(tent,\ rainy) - u(garden,\ rainy)}{u(garden,\ calm) - u(tent,\ calm)} = \frac{p_1(calm)}{p_1(rainy)} = \frac{0.3}{0.7} = \frac{3}{7}.$$

Since $u(tent,\ rainy) = u(tent,\ calm) = 0$ and $u(garden,\ calm) = 1$, we have that

$$\frac{-u(garden,\ rainy)}{1} = \frac{3}{7},$$

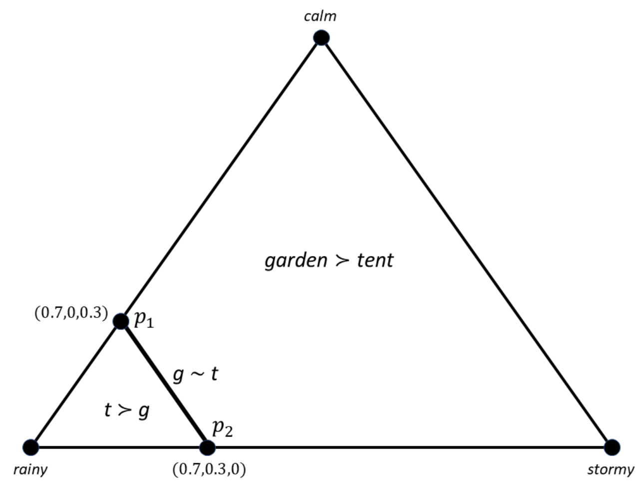

Figure 2.7.2 Question 2.4.1

which means that $u(garden,\ rainy) = -3/7$.

Next, consider the belief $p_2 = (0.7, 0.3, 0)$ in Figure 2.7.2 on the line segment between *rainy* and *stormy*. By the utility difference property,

$$\frac{u(tent,\ rainy) - u(garden,\ rainy)}{u(garden,\ stormy) - u(tent,\ stormy)} = \frac{p_2(stormy)}{p_2(rainy)} = \frac{0.3}{0.7} = \frac{3}{7}.$$

As $u(tent,\ rainy) = u(tent,\ stormy) = 0$ and $u(garden,\ rainy) = -3/7$, we have that

$$\frac{3/7}{u(garden,\ stormy)} = \frac{3}{7},$$

which means that $u(garden,\ stormy) = 1$. The utility function we obtain is thus given by

u	*rainy*	*stormy*	*calm*
tent	0	0	0
garden	$-3/7$	1	1

.

It may be verified that this utility function represents the conditional preference relation from Figure 2.7.2.

Question 2.5.1. (a) There are no preference reversals between a and b, nor between a and c. Indeed, a strictly dominates b and c.

(b) The two utility functions below both represent the conditional preference relation at hand:

u	x	y
a	2	2
b	1	0
c	0	1

v	x	y
a	3	3
b	1	0
c	0	1

.

Note that under the utility function u, the intensity by which the DM prefers a to b at x is the same as the intensity by which he prefers b to c at x. However, under the utility function v, the intensity by which the DM prefers a to b at x is twice the intensity by which he prefers b to c at x.

Question 2.6.1. Consider the utility function v given by

	rainy	*stormy*	*calm*
house	10	17	0
garden	0	7	10
tent	3	0	3
square	−10	−10	−10

.

Then, it may be verified that also this utility function represents $\succsim$. However, your choice *tent* is no longer strictly dominated by the randomized choice r under v. To see this, note that for the state *rainy,* the expected intensity by which you prefer the choice selected by r to the choice *tent* is now given by

$$(0.4) \cdot (10 - 3) + (0.4) \cdot (0 - 3) + (0.2) \cdot (-10 - 3) = -1 < 0.$$

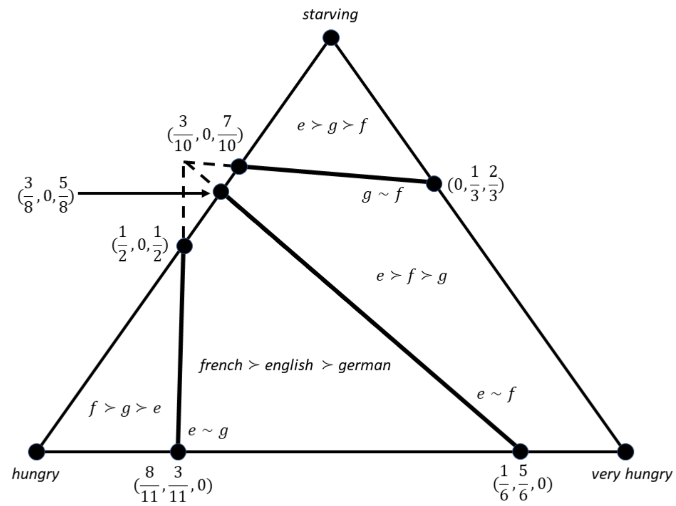

Figure 2.7.3 Conditional preference relation for Problem 2.1

Problems

Problem 2.1: The breakfast problem.

You are staying at a nice hotel with a beautiful view. It is now 8 pm in the evening, and according to the rules of the hotel you already have to order which type of breakfast you want to have for the day after. There is the choice between a huge *English breakfast* with plenty of saucages, eggs, bacon, and beans in tomato sauce, a medium sized *German breakfast* with some delicious German rolls, or a lovely but very light *French breakfast* with freshly baked croissants.

The problem is that you do not know how hungry you will be in the morning. But you are a hungry person in general, and hence you know that you will either be *starving, very hungry,* or just *hungry* in the morning. Of course, which type of breakfast you prefer will depend on how hungry you are: The huge English breakfast will be favorite if you are sufficiently hungry, but with a moderate appetite you would rather go for the small French breakfast, because you love croissants. On the other hand, you are not so fond of the German rolls.

More precisely, your preferences over the three breakfast types, for every possible belief over your appetite, are given by the conditional preference relation in Figure 2.7.3.

Here, the vector $(3/8, 0, 5/8)$ represents the belief where you think you will be hungry with probability $3/8$, you think you will be very hungry with probability 0, and you think you will be starving with probability $5/8$. Similarly for the other vectors.

(a) Use the utility design procedure to compute the unique expected utility representation u where all your utilities for the *English breakfast* are 10, and your utility for the *German breakfast* is 13 if you are just *hungry.*

(b) Explain, on the basis of Theorem 2.5.1, why the relative preference intensities are unique for this conditional preference relation.

(c) Consider the intensity by which you prefer the *French breakfast* to the *German breakfast* when (i) you are just *hungry,* and (ii) when you believe to be *very hungry* or *starving* with probability 0.5 each. What can you say about the relative proportion of these two preference intensities?

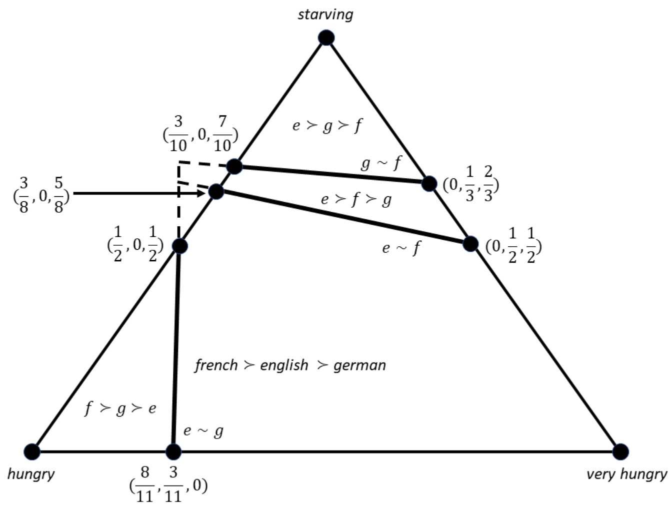

Figure 2.7.4 New conditional preference relation for Problem 2.1

(d) Suppose you believe to be *very hungry* or *starving* with probability 0.5 each. Consider the intensities by which you (i) prefer the *English breakfast* to the *French breakfast,* and (ii) prefer the *French breakfast* to the *German breakfast.* What can you say about the relative proportion of these two preference intensities?

(e) Which breakfast choices for you are rational and which are irrational? For every irrational breakfast choice, find another choice, or randomized choice, that strictly dominates it under the utility function u found in (c). Explain why it strictly dominates the irrational choice under *every* expected utility representation.

It is now one year later, and you come back to the same hotel, facing the same decision problem. Suppose that, due to a small change in your tastes, your conditional preference relation has changed. Your new conditional preference relation is given by Figure 2.7.4.

(f) Explain, without making any calculations, why this conditional preference relation does not have an expected utility representation.

Problem 2.2: Visiting Barbara.

This afternoon you will be visiting your friend Barbara, who lives at the other side of town. The problem is which type of transport to use. The possible options are going by *car,* by *bus,* by *taxi* or by *metro.* If there is no traffic jam in the city center, then you would prefer the *car* to the *bus,* the *bus* to the *taxi,* and the *taxi* to the *metro.* The reason is that you like driving, the taxi is very expensive, and you were robbed in the metro last year. However, if there is a traffic jam in the city center, then driving yourself would be a disaster, the bus would be somewhat faster because it can use the special bus tracks, the taxi driver would be even faster because he would know the quickest detours, whereas the metro would certainly get you at Barbara's place in time. Since you hate being late, you would prefer the *metro* to the *taxi,* the *taxi* to the *bus,* and the *bus* to the *car* in case there would be a traffic jam in the city center.

However, you do not know whether there will be a traffic jam in the city center or not. Your conditional preference relation can be described on the basis of the following information:

- your prefer the *car* to the *bus* precisely when you deem the probability of a traffic jam lower than 1/3;
- your prefer the *car* to the *taxi* precisely when you deem the probability of a traffic jam lower than 4/7;
- your prefer the *car* to the *metro* precisely when you deem the probability of a traffic jam lower than 1/2;
- your prefer the *bus* to the *taxi* precisely when you deem the probability of a traffic jam lower than 3/4;
- your prefer the *bus* to the *metro* precisely when you deem the probability of a traffic jam lower than 4/7;
- your prefer the *taxi* to the *metro* precisely when you deem the probability of a traffic jam lower than 1/3.

(a) Give a graphical representation of this conditional preference relation, similar to the one in Figure 2.4.7.

(b) Use the utility design procedure to compute the unique expected utility representation where all utilities for *car* are 5, and the utility for *bus* is 4 if there is no traffic jam.

(c) Explain, on the basis of Theorem 2.5.1, why the relative preference intensities are unique for this conditional preference relation.

(d) Consider the belief where you deem the probability of a traffic jam equal to 40%. What is the induced ranking of your choices for this belief? What are the relative intensities by which you prefer the number 1 choice to the number 2 choice, the number 2 choice to the number 3 choice, and the number 3 choice to the number 4 choice?

(e) Which choices are rational and which choices are not? For every irrational choice, find another choice, or a randomized choice, that strictly dominates it under the utility function found in (b). Explain why it strictly dominates the irrational choice under every utility function that represents the conditional preference relation.

It is one month later, and you again would like to visit Barbara, this time for her birthday. Since the metro ride of last month was not as bad as you thought, your preferences between *taxi* and *metro* have changed: Your prefer the *taxi* to the *metro* precisely when you deem the probability of a traffic jam lower than 1/5. All other aspects of your conditional preference relation remain unchanged.

(f) Does your new conditional preference relation have an expected utility representation? Explain your answer.

*Problem 2.3: The surprise trip.

Remember from this chapter that you celebrated your birthday with a party. You enjoyed it tremendously, but the next day you wake up with a terrible hangover. During breakfast, Barbara tells you that she has planned a surprise trip for you, in honour of your birthday, and that you have to leave within one hour. The only thing she tells you is that you will either be flying to Moscow, to Paris or to Athens. You checked the weather, and tomorrow it will be -20 degrees Celcius in Moscow, 5 degrees in Paris and 15 degrees in Athens.

There is, however, one problem: You can only take handluggage with you, which means there is only room for one outfit. Moreover, most of your clothes still have to be washed, but there is no time for this until you leave. When you stand in front of your wardrobe, there are only four outfits that

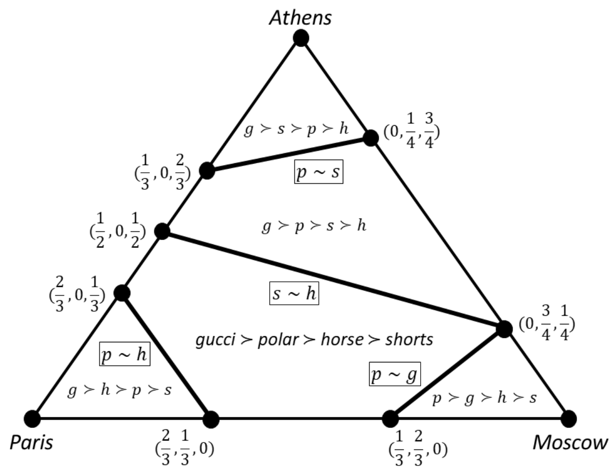

Figure 2.7.5 Conditional preference relation for Problem 2.3

you can take: a *polar suit,* a *Gucci suit,* a pair of *shorts* with Hawaï shirt, and a *horse riding costume.* Suppose that your conditional preference relation over these four outfits is given by Figure 2.7.5. The question is: Which outfit will you take?

(a) Are there outfits that are strictly dominated by other outfits? If so, which ones?

(b) We will see in (c) that there is an expected utility representation for the conditional preference relation. Explain why the relative preference intensities between *polar suit, shorts* and *horse riding costume* are unique. Can the same be guaranteed for the relative preference intensities between *all four* outfits? Explain why or why not.

Consider the signed conditional preference relation in Figure 2.7.6, which extends the conditional preference relation from Figure 2.7.5.

(c) Find the unique expected utility representation for this signed conditional preference relation where all utilities for *polar suit* are 100, and the utility for *shorts* at *Paris* is 30.
(**Hint:** Compute first the utilities for *polar suit, shorts* and *horse riding costume.* Afterwards, find the utilities for the *gucci suit.*)

Consider now the *alternative* signed conditional preference relation in Figure 2.7.7, which also extends the conditional preference relation from Figure 2.7.5. Note that the conditional preferences between *polar suit, shorts* and *horse riding costume* are the same as before, as well as the conditional preference between *polar suit* and *gucci suit.*

(d) Compute the unique utility function that represents this signed conditional preference relation where all utilities of *polar suit* are 100, and the utility for *shorts* at *Paris* is 30.
(**Hint:** Since the conditional preferences between *polar suit, shorts* and *horse riding costume* are the same as before, the utilities for these three outfits will also be the same as before.)

(e) Does the conditional preference relation from Figure 2.7.5 induce unique relative preference intensities? Explain your answer.

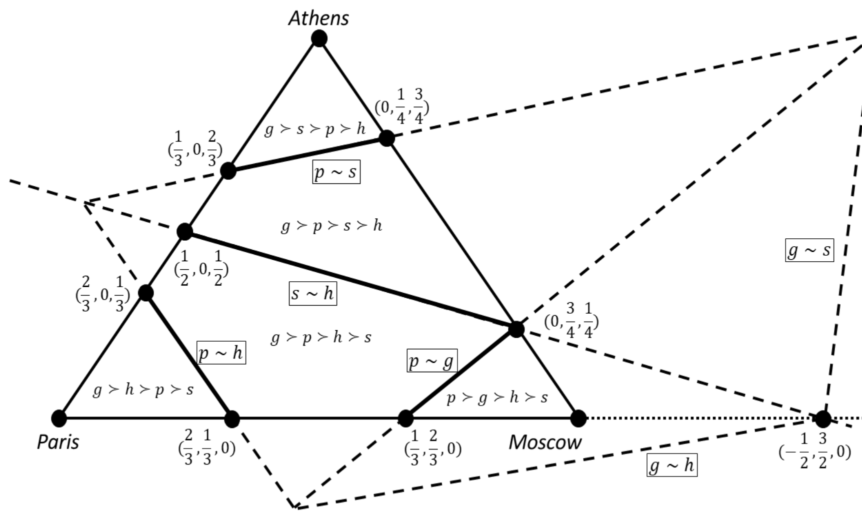

Figure 2.7.6 Signed conditional preference relation for Problem 2.3

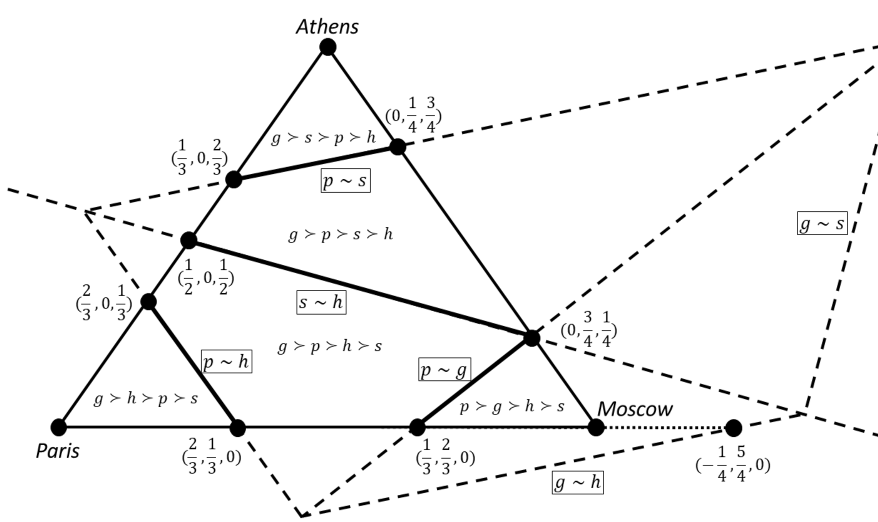

Figure 2.7.7 Alternative signed conditional preference relation for Problem 2.3

Literature

Early foundations of expected utility. Savage (1954) proposed a model of decision making under uncertainty that has become the cornerstone to modern decision theory. The primitive notions in this model are a set of *states,* representing the events about which the DM is uncertain, and a set of *consequences* (or *outcomes*), representing the possible consequences of his choices at the various states. An *act* is a mapping that assigns to every state a consequence, and Savage assumes that the DM has a preference relation over *all* possible acts.

Savage then imposes some axioms on this preference relation over acts, and shows that these axioms are both necessary and sufficient for an *expected utility representation.* Here, we say that the preference relation $\succsim$ over acts has an expected utility representation if there is a belief p over the states, and a utility function u, assigning to every outcome o some utility $u(o)$, such that the DM prefers act a to act b precisely when the expected utility of act a under the belief p and utility function u is larger than that of b. Moreover, Savage shows that under his axioms, the belief p over states will be unique, whereas the utility function u is unique up to an additive constant and a positive multiplicative constant.

Some years earlier, von Neumann and Morgenstern (1947) presented a different framework, where the DM has preferences over *lotteries* with *objective probabilities.* That is, the DM still has uncertainty about which outcome will result, but the probilities by which these outcomes will realize are *objectively given.* One could think, for instance, of lottery tickets, where the probability of winning a given price is known. Like Savage, also von Neumann and Morgenstern impose axioms on the DM's preference relation, and show that these axioms are both necessary and sufficient for an expected utility representation. Here, an expected utility representation means that there is a utility function u, assigning to every outcome o some utility $u(o)$, such that the DM prefers lottery a to lottery b precisely when the expected utility of a under the utility function u is larger than that of b. Also here, the utility function is unique up to an additive constant and a positive multiplicative constant.

Anscombe and Aumann (1963) proposed a model which is essentially a combination of the two models above. The difference with Savage lies in the definition of an act: In Anscombe and Aumann (1963), an act assigns to every state a *lottery* over consequences à la von Neumann and Morgenstern. Like Savage, also Anscome and Aumann assume that the DM has preferences over all possible acts. They impose axioms on this preference relation, and show that these axioms are necessary and sufficient for the preference relation to be representable by expected utility.

Problems with these foundations. In my view, there are at least two problems with the foundations proposed above. First, the models by Savage and Anscome-Aumann assume that the DM holds preferences over *all* possible acts – that is, over all possible mappings that assign a consequence (or lottery over consequences) to every state. Typically, however, many of these mappings will not be relevant for the decision problem at hand. Suppose, for instance, that the DM has to decide whether to take the *car* or the *bike* to work, and that he has uncertainty about whether there will be a *traffic jam* or not. Assume there are three possible outcomes: The DM can either arrive *early*, *just in time*, or *late* at work. Suppose moreover that the acts corresponding to *car* and *bike* are given by the following table:

	jam	*no jam*
car	late	early
bike	just in time	just in time

According to the models in Savage and Anscombe-Aumann, the DM should also consider, in his preference relation, the act a that would make him arrive *early* if there is a *traffic jam*, and arrive *late* if there is *no traffic jam*. Such acts, however, seem entirely artificial, and not relevant for the decision problem at hand. The same applies to the model by von Neumann and Morgenstern: Typically, there will be many lotteries over outcomes that are merely artificial mathematical constructs, and not relevant for the decision problem at hand.

The second problem is that the axioms imposed by Savage and Anscombe-Aumann lead to a *unique* belief for the DM over states. However, in many decision problems and – as we will see in later chapters – in many games, the belief of the DM may vary for several reasons. For instance, the DM may adjust his belief during his reasoning process. Or the DM may receive some new information about the states, inducing him to revise his belief. In the context of games, which will be explored in the chapters that follow, the DM must reason about the decision problems of other decision makers, and the DM will typically be uncertain about the belief that the other decision makers may hold. Given all this, it seems natural not to pin down the DM's belief, but to keep it as a variable in the model. This is precisely what we have done in this chapter, by using conditional preference relations.

Conditional preference relations. The main ingredient of our model in this chapter was the DM's *conditional preference relation,* which assigns a preference relation over choices for every possible belief that the DM may have. It differs in various dimensions from the foundations described above. First, unlike Savage and Anscombe-Aumann, we do not fix the DM's belief, but keep the belief as a variable, for the reasons explored above. Thus, the model in this chapter does not suffer from the second problem above. Moreover, the primitive notions in our model are the DM's set of *choices,* representing the options between which he must choose, and the set of *states,* representing the events about which he is uncertain, and that are relevant for the eventual consequence. By requiring that the DM, for every belief, only holds preferences over his actual *choices,* we make sure that the DM will never have to consider acts that are not part of the decision problem. Hence, we overcome the first problem outlined above.

Conditional preference relations were first introduced by Gilboa and Schmeidler (2003), although they did not call it this way. The term conditional preference relation comes from Perea (2020, 2023).

Axioms. In Chapter 2 of the online appendix to this book, we provide an axiomatic characterization for expected utility in the context of conditional preference relations. That is, we impose a list of conditions on conditional preference relations that is both necessary and sufficient for the conditional preference relation having an expected utility representation. This axiomatic characterization builds on Perea (2023). The key axioms in the characterization are the *regularity axioms, three choice linear preference intensity* and *four choice linear preference intensity.*

The regularity axioms closely resemble some of the axioms used by Gilboa and Schmeidler (2003). The axioms of *constant preference intensity* and *transitive preference sensitivity* in Jagau (2022) are similar to the axioms of *three choice linear preference intensity* and *four choice linear preference intensity*, respectively.

Similar foundations of expected utility. In Chapter 2 of the online appendix to this book we have given axiomatic characterizations of expected utility for three different scenarios: For the case of two choices, for the case where there are preference reversals for all pairs of choices, and for the general case. For the first scenario we have used the regularity axioms only, whereas for the second scenario we have supplemented these by three choice linear preference intensity and four choice linear preference intensity. For the general case we have extended the foregoing axioms to *signed* conditional preference relations, and added the axioms of transitive constant preference intensity and four choice linear preference intensity with constant preference intensity.

In the literature there are related axiomatic foundations for expected utility in the framework of conditional preference relations. These include the foundations provided by Gilboa and Schmeidler (2003), Perea (2020) and Jagau (2022).

In Gilboa and Schmeidler (2003) it is shown that if a conditional preference relation satisfies the regularity axioms together with their *diversity* axiom, then it will have an expected utility representation. The diversity axiom states that for every strict ordering of four choices or less, there must be a belief at which the DM holds precisely this ordering. However, the diversity axioms rules out many situations of interest, such as all scenarios where a choice is weakly dominated by another choice, all cases with two states and more than two choices, and all cases with three states and more than three choices. Moreover, the Gilboa-Schmeidler axioms are *sufficient,* but *not necessary,* for an expected utility representation. But, as Gilboa and Schmeidler show, their axioms are sufficient and necessary for a representation by a *diversified utility matrix,* where no row is weakly dominated by, or equivalent to, the affine combination of at most three other rows.

Perea (2020) proves that the regularity axioms, together with the axiom *existence of a uniform preference increase,* are both necessary and sufficient for an expected utility representation. The existence of a uniform preference increase states that from the conditional preference relation at hand, one should be able to increase the preference intensity between a fixed choice a and each of the other choices by a uniform amount.

In Jagau (2022) it is shown that the regularity axioms, together with the axioms of *constant preference intensity* and *transitive preference sensitivity,* are necessary and sufficient for an expected utility representation if there are preference reversals between all pairs of choices. As stated above, constant preference intensity and transitive preference sensitivity closely correspond to our axioms of three choice linear preference intensity and four choice linear preference intensity, respectively.

Unique relative preference intensities. In Theorem 2.5.1 we have shown that under certain conditions the relative preference intensities for the DM are unique. This implies that there are $|S|+1$ degrees of freedom for choosing the expected utility representation, where $|S|$ is the number of states. This result comes from Perea (2023). In contrast, Savage (1954) has shown that in his framework, the expected utility representation is unique up to an additive constant and a positive multiplicative constant. That is, there are two degrees of freedom here.

Strict dominance. Pearce (1984) shows that in two-player static games, a choice for a player is not optimal for any probabilistic belief, precisely when it is either strictly dominated by another choice, or strictly dominated by a randomized choice. In later chapters we will see that from the viewpoint of a single player, a game corresponds to a (one-person) decision problem. In that light, our Theorem 2.6.1 can thus be viewed as an extension of Pearce's result to general decision problems. In the same theorem it is shown that if a choice is strictly dominated by a randomized choice, then the randomized choice can be chosen such that it will strictly dominate the choice for *every* possible expected utility representation. This part of the theorem is due to Perea (2023).