

Decision Making under Ambiguity

13.1 INTRODUCTION

Agents must make decisions in situations characterised by uncertainty that differs both in kind and in severity. Differs in kind because they face not only factual uncertainty about the state of the world but also option uncertainty about what the consequences would be of performing one or other of the actions available to them, evaluative uncertainty about the desirability of these possible consequences and modal uncertainty about the space of relevant contingencies. Differs in severity because the quality, amount and coherence of the information that the agent has about relevant prospects can vary to a considerable degree. Mainstream Bayesian decision theory recognises some distinctions in severity (between risk and uncertainty, for instance) but measures all the different kinds of uncertainty in the same way, namely by means of a probability function defined on the set of possible states of the world.

In the first two parts of the book, I argued that such reduction of all uncertainty to factual uncertainty is not always possible or useful, and offered an alternative theory that was applicable even when it was not. Firstly, the probability measure of factual uncertainty was complemented with a desirability measure of evaluative uncertainty that explicitly incorporated dependence on both beliefs about the facts and belief-independent judgements of value, and which could be revised as these beliefs and evaluations changed. Secondly, option uncertainty was captured by a suppositional probability on prospects conditional on an intervention of some kind. These fed into a decision rule prescribing choice, from the set of available options, of the ones that maximise expected desirability gain, relative to the status quo, on the supposition of its performance. When options can be formulated as

Savage-style acts this decision rule coincides with that of maximising subjective expected utility.

This broadly Bayesian decision theory shares with its mainstream cousins the assumption that agents come to decision problems equipped with a complete set of probability and desirability judgements. It is, in other words, a decision theory suitable for a maximally opinionated agent. This implies that, if decision makers want to use such a Bayesian decision theory as a guide to their choices, then they need to reach precise judgements on at least all contingencies relevant to the decision problem they face. This view has come under considerable criticism of late, with many philosophers and economists arguing that in situations of severe uncertainty and/or irresolvable disagreement it is neither possible nor desirable for the decision maker to make precise judgements about all decision relevant contingencies nor for her to make decisions in the manner prescribed by Bayesian decision theory.

This chapter will be devoted to the examination of this contention and its implications. My starting point will be the framework of Imprecise Bayesianism developed in the previous two chapters in which an agent's uncertainty is captured by sets of pairs of probability and desirability functions and which in effect encodes the permissibility of incomplete judgement. I will start by asking how a rational but not maximally opinionated agent might reach decisions, surveying some of the many proposals that have been made in this regard. I will then focus on the question of what role considerations of caution can play in resolving decision problems, asking whether the forms of cautious decision making that are frequently observed are rational and whether they involve violations of Bayesian norms. This will give me the opportunity to take up the last of the three challenges to Imprecise Bayesianism posed in Chapter 11.

In the final chapter I will turn to a second criticism of Bayesian theory, namely that probability does not suffice to measure all the factual uncertainty that agents face. Unlike the first, this criticism extends to the Imprecise Bayesianism defended in the previous chapters. Elaboration of it will lead us to consideration of the role of confidence in judgement and to a proposal as to how confidence judgements can help to resolve some of the problems facing the Imprecise Bayesian.

13.2 REACHING A JUDGEMENT

How should a decision maker choose amongst the courses of action available to her when she lacks precise probabilities and desirabilities for the contingencies relevant to her decision? There are, broadly speaking, two possible responses to this question. Firstly, the decision maker can try and make up her mind to the degree needed to apply Bayesian decision theory, by settling on the required precise desirability and probability judgements. And, secondly, she can make use of a different decision rule from that of

maximisation of subjective expected utility; one that is much less demanding in terms of the judgemental precision it requires. In the next two sections I will put some flesh on these alternatives, without trying to settle immediately the question of which is the best route to take. Indeed, later on I will argue that different responses are applicable under different circumstances.

Until quite recently the accepted solution to the problem of decision making under severe uncertainty and/or disagreement was a version of the first response. Classical Bayesians argued that we must turn to subjective judgement for the probability and desirability values required to implement the rule of subjective expected utility maximisation, pointing out that Leonard Savage had shown that considerations of rationality require that decisions are made *as if* they maximise the decision maker's subjective expectation of benefit relative to her degrees of belief and her preferences. So precise judgements (whether explicit or not) are mandatory, on pain of irrationality, irrespective of the circumstances in which the decision is made.¹

The Classical Bayesian view is most compelling when the source of the less than maximal opinionation on the part of the agent derives from the considerations previously collected together under the label of 'Boundedness' (see Section 11.3). If, for instance, her preferences are incomplete simply because she has not given the matter much thought, then it is perfectly reasonable to expect her to put in the effort required to reach a judgement. But as a general prescription it faces two obvious challenges. The first is that under conditions of severe empirical and evaluative uncertainty and/or disagreement the decision maker may find it very difficult to arrive at a precise subjective judgement about all relevant factors. Indeed, if she is uncertain about the state space itself, she may find it impossible to do so in a non-arbitrary way. More generally, she may reasonably regard each of a range of different probability and utility judgements as equally justifiable.

The second challenge derives from the famous Ellsberg Paradox and experiments based on it, which show that under conditions of severe uncertainty many apparently rational agents seem not to conform to the dictates of Savage's theory. Those who regard this empirical evidence as normatively significant argue that it undermines the pragmatic case for expected utility maximisation and reveals a need for alternative decision rules in conditions of severe uncertainty. In reply, Bayesians argue that this experimental evidence does little to undermine the normative appeal of their theory and argue that we have no choice but to 'bite the bullet' and do the best we can to come up with reasonable probability and utility judgements

¹ As Binmore (2008) notes, Savage himself was more cautious and acknowledged that his argument was suited only to circumstances in which you could 'look before you leap' because all contingencies have been foreseen, but this qualification was largely ignored by those who followed him.

by resolving the uncertainty and/or disagreement one way or another.² John Broome, for instance, argues:

The lack of firm probabilities is not a reason to give up expected value theory. You might despair and adopt some other way of coping with uncertainty... That would be a mistake. Stick with expected value theory, since it is very well founded, and do your best with probabilities and values.

(Broome, 2012, p. 129)

Such biting of the bullet need not involve a renunciation of Imprecise Bayesianism. Contexts of enquiry impose different requirements from contexts of decision, and it is reasonable to hold the view that the judgements that an agent takes as the basis for action may be more precise than those she forms on the basis of the evidence she holds. In enquiry it is the Jamesian imperative to avoid error and to keep an open mind that takes precedence; in decision making it is the Jamesian imperative to seek truth and to form an opinion that does. So long as the precise opinions formed for the purposes of making a decision can subsequently be suspended if there is opportunity for further enquiry, there is no reason to fear spells of pragmatic dogmatism (hence the importance to Imprecise Bayesianism of the rule of opinion withdrawal described in the previous chapter).

Clearly, though, this imperative to form opinions when they are called for needs to be backed up with some advice on how this might be done when the decision maker lacks information about relevant contingencies or is divided in her evaluations. If she has the time, she can seek further information or deliberate further in the hope that this will help to settle matters. This would normally require postponement of the decision, a possibility examined in the next section, so I shall set it aside for the moment. If postponement is ruled out there are two (not mutually exclusive) strategies she can pursue: she can try and identify the ‘best’ opinion by applying additional (non-evidential) considerations; or she can try and form an aggregate of the permissible opinions that is, in some sense, the best compromise between them. Let’s look at some examples of each.

13.2.1 Picking the ‘Best’

We have already encountered a salient version of the first strategy. When the evidence does not fully discriminate between various hypotheses, objective Bayesians look to the Principle of Indifference to determine a unique probability assignment. Recall that application of this principle to a set $O = \{O_i\}$ of mutually exclusive and exhaustive outcomes leads, in the absence of any information distinguishing these outcomes, to an assignment of equal probability to each. But what about when we have partial information about these outcomes? Then, it is argued, we should pick a probability

² See, for instance, Nabil Al-Najjar and Jonathan Weinstein (2009).

assignment consistent with this information that departs minimally from an equal assignment to all outcomes. On a natural metric for minimal departure, this yields the rule of Maximum Entropy (MaxEnt).³ As its name suggests, this rule picks the member P of the set of probability functions consistent with the evidence the agent holds which maximises entropy relative to the set of outcomes – i.e. it minimises the measure

$$H(P) = \sum_i P(O_i) \cdot \log(P(O_i))$$

By departing minimally from the equal assignment, we avoid giving any more probability to outcomes than the evidence requires us to. And, in that sense, we thereby adopt the most equivocal set of degrees of belief that we can, given the evidence we hold. But what reasons do we have to be equivocal in this sense? Jon Williamson (2007b) argues that the adoption of equivocal degrees of belief leads, on average, to more cautious decision making and hence that caution furnishes pragmatic grounds for MaxEnt. Suppose I don't know the bias on a coin that is to be tossed and must bet on how it lands. An epistemic 'equivocator' will bet cautiously in the sense of refusing bets that pay \$1 on heads that cost more than 50c. In contrast, the epistemically 'reckless', who adopts a probability of one for heads, will accept a bet at any price up to \$1 for the bet. On the other hand, however, while Reckless will not sell such a bet at any price, Equivocator will willingly sell at 50c.

Is there anything that we can say in favour of or against one or other of these betting decisions? Nothing much at all of substance, it would seem. We don't know how the coin will land; indeed, we don't even know what its chance of landing heads is. So, we not only don't know who will do better as a result of her choices, we can't even say who can be expected to do better. If Reckless pays \$1 and loses then she loses big; bigger anyway than Equivocator who pays only 50c. So equivocation can help to minimise losses. But it also minimises gains. For Equivocator will forgo opportunities to bet and some of these will pay out for Reckless. Only if losses matter more to the agent than gains is there anything to be said in favour of caution. But then the attraction of equivocation cannot be a purely epistemic matter, nor can it be completely general. So the claim that equivocation is objectively required or rationally mandatory seems completely without foundation.

When the dust settles we are left with little more than the original thought behind the Principle of Indifference: that absence of evidential reasons for differential probability assignments is a reason to make equal ones. That this reason is not itself an evidential reason should not be held against MaxEnt since our problem is precisely to pick a probability when evidential reasons give out. And at least it does provide the subjectivist Impreciser with a way to settle on a precise opinion for the purposes of making a decision.

³ See Williamson (2007a), Landes & Williamson (2013), Paris (2006) and Jaynes (1968, 2003).

There are dangers here for the subjectivist, however, and she should consume with moderation. For one thing, the Principle of Indifference is notoriously sensitive to the choice of description of outcomes, yielding different prescriptions depending on the level of refinement of the problem; a difficulty that carries over to the MaxEnt rule. And, for another, application of MaxEnt can conflict with Bayesian conditionalisation (see Williamson, 2011). Somewhat ironically, however, neither need be too much of problem for the subjectivist seeking only to apply MaxEnt for decision purposes. For in many decision-making contexts, and in particular those characterised by what I previously called Grade 2 uncertainty, what is needed is a probability assignment for the decision problem *as it is formulated*. A subjectivist wishing to avail herself of MaxEnt will represent her decision problem in a way that she judges is appropriate for application of the rule – i.e. where the symmetries that she takes to be present in the situation she finds herself in are captured in the descriptions of the outcomes. Furthermore, in using MaxEnt to determine a probability for the decision at hand she need not commit herself to holding onto these probabilities in the future and can perfectly well opt to suspend opinion again after the decision has been made and implemented. So conflict with Bayesian norms of belief revision can be avoided.

13.2.2 Aggregating

The second strategy that can be pursued is to reach an opinion by aggregating all those opinions that the agent regards as permissible. In essence, the idea is to exploit an analogy between a group agent and an individual agent with multiple avatars in disagreement on the question of what opinion to adopt, and then to draw on the large literature in social choice theory to provide rules for fashioning precise aggregate judgements.⁴ Two classes of aggregation rules are particularly salient in this literature. Voting rules select the opinion with the greatest support by counting the number of voters (in this context, avatars) endorsing it, perhaps weighting them on the basis of other considerations, such as how competent they are or how affected they are by the decision. Averaging rules, such as ‘splitting the difference’ and linear averaging, on the other hand, select opinions that are the best compromises between the individual ones. The Principle of Indifference makes an appearance here too. When there is no reason for favouring one opinion over the others in virtue of who holds it or its

⁴ Methods for aggregating different kinds of opinion have been extensively studied. The problem of aggregating probabilities has received most attention from statisticians (see Genest & Zidek, 1986, for a survey), while Social Choice theory has mainly focused on the problem aggregating preferences and/or utilities (see Sen, 1970, for a classic discussion), but there is some work on the joint aggregation of probabilities and utilities (see, for instance, Mongin, 1995, and Bradley, 2005a). Finally the theory of judgement aggregation, developed by Christian List, Franz Dietrich and others, tackles the problem of aggregation in a very general way (see List & Puppe, 2009, for a recent survey).

content, the principle dictates treating each equally: rules such as majority rule and equal weighted averaging respect this dictum. But the literature recognises a great variety of contexts in which considerations favour either particular opinion holders or particular propositions and provides aggregation rules appropriate to them.

Some perspective on the scope and limits of such rules can be gained by looking in more detail at one of the most widely endorsed proposals, namely that an agent should form an aggregate of a set of probability judgements by taking a weighted average of them. Formally, given a set $C = \{P_i\}$ of probability functions defined on a common domain of propositions, a linear average, P_0 , of these probabilities is obtained by setting, for some set of corresponding weights $\{w_i\}$ such that $w_i > 0$ and $\sum_i w_i = 1$:⁵

$$\text{Linear Averaging } P_0 = \sum_i w_i \cdot P_i$$

There are a variety of possible interpretations of the probabilities and the weights on them occurring in this formula. One rather salient one, particularly relevant to the context of uncertainty, treats the P_i as the various candidate hypotheses as to the true probabilities or objective chances of the prospects in their common domain and the weights as second-order probabilities on an extended domain containing chance propositions. So interpreted, Linear Averaging is simply an implication of the Principal Principle. And the proposal it supports, namely to adopt the expected chances of prospects as one's aggregate degrees of belief for them, amounts to what might be called 'second-order Probabilism', since it enriches the standard probabilistic framework in a way which allows for rationality constraints rooted in beliefs about objective probabilities.

The obvious problem with this proposal in this context is that it is hard to see how agents who lack the information necessary to form first-order probability judgements would nonetheless be able to form second-order ones, in particular to assign probabilities to hypotheses about what the chances are. Indeed, second-order Probabilism is more demanding cognitively than the simple first-order version, which requires no recognition on the part of the agent of objective chances. So this interpretation doesn't usefully apply in contexts of severe uncertainty insofar as we are interested in providing agents with guidance as to how to deliberate about their uncertainty.

A second interpretation, more appropriate to contexts of disagreement, treats the P_i occurring in Linear Averaging as the probabilistic judgements of different experts and the weights as some measure of either the experts' competence or reliability or the confidence the agent has in them (which is to be distinguished from confidence in a belief or a probability judgement).

⁵ The stipulation of strictly positive weights reflects the assumption that no probability function not deserving of a positive weight should be in the permissible set in the first place.

Lehrer & Wagner (1981), for instance, promote this as a method for forming opinions in cases of disagreement amongst experts, including ones in which the ‘experts’ are simply the agent’s epistemic peers. A variant treats the P_i as the outputs of different models or lines of enquiry, with the weights once again measuring either the reliability of the methods they employ or the confidence that the agent has in them (see, for instance, Gärdenfors, 1988). On both, the aggregate probability P_0 can be construed as a confidence-weighted average of the probability judgements that the agent regards as worthy of consideration.

Contrary, however, to the claim of Lehrer (1976, 1983) that linear averaging is the uniquely rational way of forming one’s beliefs in the face of disagreement amongst experts, it has a number of significant weaknesses. I’ll mention two here. The first is that this rule is insensitive to whether the opinions expressed by different individuals on the same proposition are independent or not. But compare a situation in which two scientists conduct separate experiments to try and settle some question with one in which they conduct a single experiment together. Suppose that in both cases the scientists report that as a result of their experiments they consider X to be highly probable. In the former case, we would want to raise our own probability for X quite considerably because of the convergence of their expert testimony. In the latter case too we would want to raise our probability for X , but less so, because their joint testimony in favour of X is based on the same information. To revise once in the light of the testimony of the first scientist and then again in the light of that of the second would in effect be to update twice on the same evidence, akin to an individual scientist conditioning twice on the same experimental result.

A second problem is that the confidence weights that this rule places on the different experts or models are proposition-*independent*. Consider a simple case in which we consult two scientists with different domains of expertise, one being an oceanographer and the other a meteorologist. It would be natural to have more confidence in what the former says about sea temperature but more confidence in what the latter says about cloud formation. But Linear Averaging requires us to assign the same weight to the probabilistic judgements of the oceanographer and the meteorologist on the second question as we do to the first. So it asks us to apply confidence considerations in an unsatisfactory way.

Could we not avoid the problem by employing proposition-*dependent* confidence weights? Unfortunately not, for if we do so then we will land up with incoherent degrees of belief (see Bradley, 2007b). Consider the following example. Suppose that you know that Anne has observed that A is true while Bob has observed that B is true. Suppose, furthermore, that they report the probabilistic degrees of belief across the partition $\pi = \{AB, A\neg B, \neg AB, \neg A\neg B\}$, as displayed in Table 13.1. Now Anne and Bob’s observations make them maximally reliable, respectively, on the question of whether or not A is true and whether or not B is true (supposing absence of observational error). So we should simply adopt their reported beliefs as our own, leaving us with degree

TABLE 13.1. *Linear Averaging*

	AB	$A \neg B$	$\neg AB$	$\neg A \neg B$
<i>Anne</i>	0.1	0.9	0	0
<i>Bob</i>	0.1	0	0.9	0
Linear average	0.1	$0.9a$	$0.9b$	0

of belief 1 in AB and 0 in all the other propositions. But this goes against the recommendations of Linear Averaging, which, irrespective of the weights a and b assigned to Anne and Bob, will yield a probability of 0.1 for the proposition AB and 0 for $\neg A \neg B$.

We can put the problem slightly differently. If you form your beliefs by averaging Ann and Bob’s opinions on the four element partition π using weights a and b for Anne and Bob, respectively, then your degrees of belief on the partitions $\{A, \neg A\}$ and $\{B, \neg B\}$ must, on pain of probabilistic incoherence, be linear averages of Anne’s and Bob’s degrees of belief obtained by applying weights a and b . But Anne and Bob have different competencies over these partitions, so these weights cannot be adequate to both.

These problems are not peculiar to linear averaging; any of the usual averaging rules found in the literature will face similar ones. Indeed, the root of the problem, it seems to me, lies in the way that confidence considerations are applied by such rules, namely as weights on experts or models rather than on the probabilistic judgements that are supported by what these experts or models say. But it may be that we are simply asking too much from these techniques by making them live up to epistemic standards appropriate to opinion formation rules that are designed to be sensitive to the evidence. For our problem is precisely how to form an opinion when such evidence gives out. In which case, the rationale for the adoption of these techniques may simply lie in the fact that they deliver a consistent solution to this problem.

On this line of reasoning, it might be reasonable to seek pragmatic grounds for the assignment of weights to experts. But letting pragmatic factors shape belief formation risks putting the cart before the horse, as we usually want our decisions to be guided by our beliefs rather than the other way around. Arguably, therefore, the appropriate point to apply pragmatic considerations is at the moment of choice, rather than during attitude formation.⁶ So let’s turn to the second strategy, of leaving attitudes imprecise when circumstances do not warrant greater precision and applying an alternative decision rule.

⁶ There are grey areas, I think. In group decision making it is often impossible to reach a decision on what to do without getting agreement on the reasons that ground that choice. In which case, pragmatic considerations are likely to slip into the phase of attitude formation.

13.3 ALTERNATIVE DECISION RULES

A decision maker who is unable or unwilling to form precise probability and desirability judgements on all prospects relevant to the decision problem she faces cannot, of course, choose in accordance with expected utility maximisation. But she might instead apply a different decision rule, one that is tailored to her state of severe uncertainty or conflict. A great many different proposals for such rules exist in the literature, involving more or less radical departures from Bayesian theory and varying in the informational demands they make. Our focus will be on rules that take as inputs the set of expected utilities associated with an act that characterise the decision maker's uncertainty, organising them in terms of the additional considerations they appeal to in order to settle matters.

More formally, let us suppose that the agent is in decision situation $\mathcal{D} = \langle O, S \rangle$, with $O = \{f, g, \dots, h\}$ the set of actions available to her and $S = \langle \Omega, \succeq, \succsim \rangle$ her judgemental state with, as before, \succeq and \succsim , respectively, being her credibility and preference relations on the Boolean algebra of prospects $\Omega = \langle X, \models \rangle$. It will not matter to the discussion here whether we think of actions in the manner of Savage, as functions from a set of states (a partition of X in terms of the features of the world that are causally independent of the actions) to consequences (a partition of X in terms of the features of the world that matter to the agent), or in terms of the partitioning indicative conditionals that pick them out.

Let $\mathcal{J} = \langle \Omega, \mathcal{A} = \{A_i\} \rangle$ be the set of avatars of the agent determined by the relations \succeq and \succsim , with each avatar being a pair of probability and desirability functions that jointly represents them. For any action f , let $\mathbb{E}_i(f)$ be the expected utility of f according to avatar i – i.e. its expected utility calculated by applying the pair of probability and desirability functions constituting that avatar.

The problem the agent faces is to settle on a choice of action on the basis of the set of expected utilities associated with each of her options and any other considerations that she can apply. One principle of choice commands universal assent: that, if all the agent's avatars assign higher expected utility to one action than another, then the latter should never be chosen when the former is available. More formally, if C is a choice function on decision situations $\mathcal{D} = \langle O, S \rangle$, then:

Unanimity If $f, g \in O$ and for all $A_i \in \mathcal{A}$, $\mathbb{E}_i(f) \geq \mathbb{E}_i(g)$, then

$$g \in C(\mathcal{D}) \implies f \in C(\mathcal{D})$$

Unanimity can be strengthened a bit by adding a second clause to the effect that if, in addition, for some $A_{i^*} \in \mathcal{A}$, $\mathbb{E}_{i^*}(f) > \mathbb{E}_{i^*}(g)$, then $g \notin C(\mathcal{D})$. But even so strengthened the Unanimity principle is unlikely to resolve the agent's decision problem in a significant number of contexts. So further considerations will have to be brought to bear on the problem in order to resolve it.

There are four such considerations that are particularly salient: *caution*, *confidence*, *robustness* and *flexibility*. To study how they can be used to settle decisions, let \geq be a ranking of options in terms of their choice-worthiness and which determines what the agent can permissibly choose. Formally, if C is a choice function on decision situations $\mathcal{D} = \langle O, S \rangle$, then for all $f, g \in O$

Choice-Worthiness Ranking $f \geq g \Leftrightarrow [g \in C(\mathcal{D}) \Rightarrow f \in C(\mathcal{D})]$.

Unanimity implies that \geq contains the intersection of the avatars' preference orderings. Let us now look at how additional considerations can be used to complete it.

13.3.1 Caution

When a decision maker regards a range of probabilities and/or desirabilities as reasonable, she may wish to be cautious in her decisions by giving more weight to the 'downside' risks – the possible negative consequences of a choice of action – and less to the 'upside' chances. Someone who is cautious in this sense will tend to hedge against risks by choosing actions with less variance in their expected outcomes. Hedging seems particularly compelling when the costs and benefits of an action in each state of the world accrue differently to different individuals, for in this case reducing the variance can serve the goal of treating individuals more equally. But, in general, it has the advantage of assuring the agent that her expected losses will not exceed some amount.

A salient decision rule encoding such caution is the maximin-EU (MMEU) rule, which recommends picking the action with the greatest minimum expected utility. In its usual formulation the expected utilities are determined relative to a fixed single utility and a set of probabilities, but it is very naturally generalised to the case in which both utilities and probabilities are imprecise. This more general MMEU rule says

MMEU $f \geq g \Leftrightarrow \min[\mathbb{E}_i(f)] \geq \min[\mathbb{E}_i(g)]$

MMEU and near variants have been advocated by a number of philosophers and economists, including Levi (1990), Peter Gärdenfors and Nils-Eric Sahlin (1982) and Itzhak Gilboa and David Schmeidler (1989), and the latter have provided an elegant representation theorem for it. Arguably, however, the rule is much too cautious, paying no attention at all to the full spread of possible expected utilities, though the force of this criticism somewhat depends on how the range of probabilities and associated expected utilities is determined in any given decision situation.

These problems can be avoided to some extent by adopting one of the rules for decision making under ambiguity, which draw on further information about the set of expected utilities determined by the agent's imprecise beliefs. Daniel Ellsberg (1961), for instance, proposes maximising a weighted average

of the minimum and mean expected utility, where the relative weights on the minimum and mean can be thought of as either reflecting the decision maker's pessimism or her degree of caution. This rule yields much the same prescriptions as maximisation of a weighted average of the maximum and minimum expected utility (often called the α -MEU or Hurwicz rule). Formally, this latter rule dictates that, for some $\alpha \in [0, 1]$, canonically taken to be a measure of the agent's degree of pessimism or caution and assumed to be greater than 0.5:

α -MEU $f \geq g \Leftrightarrow$

$$\alpha \min_i [\mathbb{E}_i(f)] + (1 - \alpha) \max_i [\mathbb{E}_i(f)] \geq \alpha \min_i [\mathbb{E}_i(g)] + (1 - \alpha) \max_i [\mathbb{E}_i(g)]$$

The α -MEU rule has been defended by Binmore (2008) and axiomatically characterised by Paolo Ghirardato, Fabio Maccheroni and Massimo Marinacci, 2004. Like Ellsberg's proposal, it generalises MMEU by allowing decision makers with the same imprecise beliefs to differ in the degree of caution that they display in their choices.

A question that all such rules must address is the specification of the set of probabilities that the expected utilities are based on. When the evidence does not determine a single probability then the Bayesian insistence on a single probability seems too extreme. But, if all probabilities *consistent* with the evidence are included, then it is likely to determine very wide probability intervals for decision-relevant contingencies. In the case of the MMEU rule, with $\alpha > 0.5$, this will tend to lead to very cautious decision making; in all cases the extremes of the probability intervals have considerable influence on the choice of action. A natural thought is that the set should determine intervals that are sufficiently broad that the decision maker is confident that the 'true' probabilities lie within them or that they contain all reasonable values. For instance, if the source of these probabilities is the opinions of others, the decision maker does not need to consider every possible opinion consistent with the evidence, only those that they have some confidence in. But how confident do they need to be? We return to this question later, once we have discussed the notion of confidence in more detail.

13.3.2 Confidence

A second set of alternative rules draws on considerations of confidence and/or reliability. The thought here is that, even if you do not know what the 'true' expected utility of an action is, you can be more or less confident about the various candidate estimates. For instance, when the estimates derive from different models or experts, the decision maker may regard some models as better corroborated by available evidence than others or some experts as more reliable than others in their judgements. In these cases it is reasonable, *ceteris*

paribus, to favour actions of which you are more confident that they will have beneficial consequences. One way of doing this is to weight each of the expected utilities associated with an action in accordance with how confident you are about the judgements supporting them and then choose the action with the maximum confidence-weighted expected utility. Formally, given a set of weights $\{\alpha_i\}$ such that $\alpha_i > 0$ and $\sum_i \alpha_i = 1$, this rule counsels choice in accordance with

$$\text{CWEU } f \geq g \Leftrightarrow \sum_i \alpha_i \cdot \mathbb{E}_i(f) \geq \sum_i \alpha_i \cdot \mathbb{E}_i(g)$$

Note the ‘kinship’ of CWEU with Linear Averaging, the rule for forming aggregate probability judgements that we looked at in the previous section. Indeed, when the agent has complete preferences for consequences so that her avatars disagree fundamentally only in their beliefs, then the ranking over acts induced by CWEU is just the same as that induced by expected utility maximisation relative to a linear average of the avatars’ probabilities.⁷

There is room here too for different interpretations of the confidence weights occurring in the CWEU equation. In some cases they can be construed as measures of the reliability of the expert or model that are their source. In other cases they can be construed as second-order probabilities; for instance, the probability that the expectation \mathbb{E}_i is the best one to use in evaluating the action. In this case, CWEU becomes a form of ‘second-order’ Bayesianism according to which the value of action is determined by the subjective expectation of its ‘true’ expected utility. As such it does little to address the problem of decision making under severe uncertainty, since it returns us to the problem of how to form precise second-order beliefs. If combined with one of the techniques for forming a judgement described in the previous section, some progress can be made, however. For instance, if there are no grounds for greater confidence in any one of the expected utility judgements than another, appeal might be made to the Principle of Indifference to motivate the assignment of equal confidence weights to the agent’s avatars. In this case, CWEU reduces to a well-known rule for decision making under conditions of ignorance: maximisation of mean expected utility. Formally:

$$\text{MaxMean } f \geq g \Leftrightarrow \sum_{i=1}^n \frac{\mathbb{E}_i(f)}{n} \geq \sum_{i=1}^n \frac{\mathbb{E}_i(g)}{n}$$

Such second-order Bayesianism – and, indeed, simple maximisation of CWEU under any interpretation of the confidence weights – leaves no room for the kind of caution considered before. But a close variant of it, the ‘smooth

⁷ There are hidden dangers here, though. When the agent’s avatars disagree in both their beliefs and their desires, then the linear average of their expected utilities may not cohere with the agent’s incomplete preferences. For discussions of these aggregation problems, see Mongin (1995) and Bradley (2005a)

ambiguity' model of Peter Klibanoff, Massimo Marinacci and Sujoy Mukerji (2005), allows for an aversion to wide spreads of expected utilities, by valuing actions in terms of a linear average of a concave transformation of their expected utilities, rather than in terms of the expected utilities themselves, where this transformation reflects the agent's degree of aversion to the spread. Formally, let $\phi : \Re \rightarrow \Re$ be such a concave mapping on the real numbers. Then, according to the smooth ambiguity model of Klibanoff, Marinacci & Mukerji,

$$\text{SAM } f \geq g \Leftrightarrow \sum_i \alpha_i \cdot \phi(\mathbb{E}_i(f)) \geq \sum_i \alpha_i \cdot \phi(\mathbb{E}_i(g))$$

A second model, due to Alain Chateauneuf and José Faro (2009), combines consideration of confidence and caution in a quite different way. They postulate a confidence threshold for determining the set of probabilities relative to which the decision maker applies the maximin-EU decision rule. In doing so they partially resolve the problem of the determination of the set C of priors, though they do not say anything about what level of confidence should be required. But considerations of confidence can be used even when precise confidence weights cannot be provided, though they then need to be supplemented with other considerations (such as caution). Peter Gärdenfors and Nils-Eric Sahlin (1982), for instance, suggest simply excluding from consideration any estimates that fall below a reliability threshold and then picking cautiously from the remainder. Similarly, Brian Hill (2013) uses an ordinal measure of confidence that allows for stake-sensitive thresholds that can be combined with other considerations. We will return to Hill's model later.

13.3.3 Robustness

A third consideration that can be appealed to is the robustness of the decision rationale.⁸ The basic thought here is that the decision maker should work out which dimensions of uncertainty make the most difference to the outcomes of her decisions and then choose actions that do sufficiently well for a reasonable range of values on these dimensions. Actions chosen on this basis will usually be 'regret-free' in the sense that, even if they do not always turn out to be optimal, they are likely not to turn out to have been a bad choice.

What counts as a reasonable range of values? Most approaches that appeal to robustness assume that a best estimate or preferred model is available and then consider small deviations away from this estimate or small changes in the model parameter values (see, for instance, Ben-Haim, 2006, and Hansen & Sargent, 2001). A robust action is one that can be expected to have beneficial consequences relative not just to the best estimates of the values

⁸ See especially Nehring (2009a), who uses the criterion of the robustness of a decision rationale in a more far-reaching way than examined here.

of relevant variables but also to a class of estimates that deviate from the 'best' one to some degree. The wider the class in question, the more robust the action. When the option that maximises expected utility, relative to the best estimate, is robust in this sense, then one gets an extra reason to choose it. But sometimes the expected utility-maximising option may be less robust than alternatives that are nonetheless satisfactory in terms of their expected utility. Then some trade-off between the two considerations, expected utility and robustness, must be made in order to resolve the question of what to choose.

13.3.4 Flexibility

The final consideration that can be appealed to is flexibility. In some contexts, an option that is available to decision makers is to delay all or part of the decision until more information is available or some of the disagreement is resolved through deliberation. The basic motive for delaying a decision is to maintain the ability to respond flexibly to contingencies that arise. Suppose, for instance, that a choice must be made between building a cheap, but low, sea wall or a high, but expensive, one, and that the relative desirability of these two courses of action depends on unknown factors, such as the extent to which sea levels will rise. In this case it would be sensible to consider building a low wall first but leave open the possibility of raising it in the future. If this can be done at no additional cost, then it is clearly the best option: at worst, no new information is acquired by the time the low wall is completed and you are in much the same situation as you started; at best, you are able to make the optimal choice at the later time. Typically, of course, flexibility comes at a cost and some judgement must be made as to whether the cost is worth bearing (and this decision may be no easier to make than the initial one). So the extent of the benefit that can be extracted by pursuing this strategy will depend on the possibility of keeping these costs down by breaking the original decision problem down into relatively autonomous, subsidiary decision problems that can be settled sequentially.

A preference for flexibility is reasonable even under conditions of normal uncertainty, when the decision maker has precise probabilities for the future contingencies that determine how beneficial a course of action is. Indeed, there are well-established models of dynamic decision making that exploit this fact (see Kreps & Porteus, 1978, and Arrow & Fisher, 1974, for instance). But the central principle at stake here, namely that, *ceteris paribus*, you should prefer actions that leave more options open to those that restrict them, can do even more important work in conditions of severe uncertainty. For, if you are unable to determine all the consequences of your possible actions or if you are unable to predict what value you will attach to these consequences at their time of realisation, then you have a strong incentive to avoid making irreversible

TABLE 13.2. *The Ellsberg Paradox*

	<i>Red</i>	<i>Black</i>	<i>Yellow</i>
L_1	\$100	\$0	\$0
L_2	\$0	\$100	\$0
L_3	\$100	\$100	\$0
L_4	\$0	\$100	\$100

commitments too early on. Finally, when you find yourself in a situation of conscious unawareness, when you are unsure about true state space and believe that you may not be aware of all relevant contingencies, the ability to respond flexibly to changes in circumstances becomes crucial.

13.4 CAUTIOUS DECISION MAKING

Drawing on additional considerations such as caution, confidence, robustness or flexibility yields a wide range of alternative decision rules to that of maximisation of expected utility, each with pretensions of being appropriate in at least some circumstances of severe uncertainty and/or disagreement. Exhaustive examination of each of these considerations and associated proposed rules is beyond the scope of this book and I will focus on the role of just two of them. The rest of this chapter will be devoted to an assessment of the role of caution in decision making, with particular attention to the question of the rationality of cautious attitudes. The next chapter will look at the role of confidence.

Much of the current debate about what decision rules are appropriate to conditions of severe uncertainty has been driven by the desire to explain the pattern of preferences frequently observed in the Ellsberg Paradox (exhibited again, for convenience, in Table 13.2) and attributed by Ellsberg to an attitude that has come to be called ambiguity aversion. Behaviour consistent with the kind of aversion to ambiguity postulated by Ellsberg has been established in numerous experiments involving set-ups similar to his.⁹ And his characterisation of it, as a type of cautionary attitude that leads decision makers to prefer actions with known, or less uncertain, chances of reaching their goals, to those with more uncertain chances, has also found widespread acceptance. On the other hand, there is considerable disagreement about how to model ambiguity aversion and on the question of whether it is rational; in particular, whether it is compatible with the Bayesian theory of rationality.

⁹ See Wakker (2010) and Trautmann & Van De Kuilen (2016) for a review of the literature.

Many decision theorists take the Ellsberg Paradox as evidence that agents do not have precise probabilities for draws of the black or yellow balls from the ‘ambiguous’ urn and then use this fact together with one of the decision rules for agents with imprecise probabilistic beliefs previously introduced to explain the Ellsberg preferences. The Maximin EU model, which prescribes choice of the alternative that maximises the minimum expected utility, affords one such explanation. If we set the utility of \$100 to 1, for instance, then the minimum expected utility of L_1 is $\frac{1}{3}$, of L_2 is 0, of L_3 is $\frac{1}{3}$ and of L_4 is $\frac{2}{3}$. So agents who employ MMEU will have the characteristic Ellsberg preferences: $L_4 \succ L_3$ and $L_1 \succ L_2$. They are not by any means the only ones, though. Agents who employ the α -MEU¹⁰ – or, indeed, *any* other of the cautious rules that were presented in the previous section – can also display these preferences and, more generally, some form of ambiguity aversion.

Although these models differ in various ways, they all take it as given that the Ellsberg preferences and, more generally, ambiguity aversion are inconsistent not only with Savage’s Sure-Thing Principle but also more generally with the view that individuals base their decisions on precise probabilities for the contingencies upon which the consequences of their choice depend. I have already argued (in Chapter 9) that this conclusion cannot be drawn from the Ellsberg Paradox and that the Ellsberg preferences are perfectly compatible both with Precise Probabilism and with expected utility maximisation, provided that agents do not value chances linearly. In this section I will develop this discussion by addressing three questions in more detail. (1) What kind of an attitude is ambiguity aversion? (2) Is ambiguity aversion rational? (3) Is ambiguity aversion consistent with Bayesian principles?

The first step to addressing these questions is to define ambiguity aversion more precisely. To do so, I will follow the characterisation of it given by Schmeidler (1989) and Itzhak Gilboa and David Schmeidler (1989) as a preference for hedging or randomisation, using a propositional variant of the framework developed by Anscombe & Aumann (1963) in which most of the recent debate on decision making under ambiguity has been conducted. In the Anscombe and Aumann (A–A) framework, actions are represented as functions from the set of states of the world to von Neumann and Morgenstern lotteries. So an action is something that determines for each state of the world an objective chance distribution over the set of prizes. Formally, let $S = \{S_1, S_2, \dots, S_m\}$ be the set of states of the world, which we continue to consider to be a partition of the set of prospects in terms of all combinations of features of the world that are causally independent of the agent’s actions but relevant to determination of the outcome of choosing a particular one. As before, let the $\bigwedge_{i=1}^n (Ch(X_i) = x_i)$ be lottery propositions: conjunctions of the propositions $Ch(X_1) = x_1$, $Ch(X_2) = x_2$, ..., and $Ch(X_n) = x_n$ that specify the chances of obtaining each of the ‘prizes’ represented by the propositions X_i .

¹⁰ See definition in Section 13.3.1.

Let Δ be the set of all such lottery propositions. As lottery propositions serve as consequences in the A–A framework, for them an action is just a function from \mathcal{S} to Δ . Let \mathcal{F} be the set of all such A–A actions and \mathcal{H} be the subset of them consisting of actions with constant lottery consequences. For any $f, g \in \mathcal{F}$ and $\alpha \in [0, 1]$, let the α -mixture of f and g , denoted by $\alpha f + (1 - \alpha)g$, be an act whose consequence in each state of the world, S , is defined by

$$(\alpha f + (1 - \alpha)g)(S) = \alpha f(S) + (1 - \alpha)g(S)$$

In other words, an α -mixture of f and g is an act whose consequences lie between those of f and g in each state of the world. Note that $\alpha f + (1 - \alpha)g$ itself belongs to \mathcal{F} . It follows that \mathcal{F} is closed under mixing.

Let \succsim be a complete and transitive preference relation on \mathcal{F} . Anscombe and Aumann make two contentious assumptions about \succsim . Firstly, they assume that preferences for lotteries are state-independent. This assumption is no more plausible as a general principle in this framework than in Savage's, but since it is not at the centre of the questions we need to consider let's just accept it for convenience. Secondly, they assume that preferences satisfy a strong separability condition on mixtures of acts which says that, if any two acts are mixed with a third one, then this does not affect the preference ranking of them. Formally, for all $f, g, h \in \mathcal{F}$ and $\lambda \in [0, 1]$:

A–A Independence $\lambda f + (1 - \lambda)h \succsim \lambda g + (1 - \lambda)h \Leftrightarrow f \succsim g$

A–A Independence is a *very* powerful axiom. Not only does it imply both Savage's Sure-Thing principle and the von Neumann and Morgenstern Independence axiom, but in fact it implies much more than the conjunction of them. For example, suppose that you must choose between betting for or against a coin landing heads that is known to be either two-headed or two-tailed (the ambiguous options) or betting on a coin landing heads that is known to be fair (the risky option). These three options are represented in Table 13.3, respectively as the bets H, T and F, with the cell entries specifying the chances of winning the bet in each of the unknown states of the world. The bet H (T) on the ambiguous coin landing heads (tails) yields a chance one of the prize in the event that the coin is two-headed (two-tailed) and a chance zero of the prize in the event that it is two-tailed (two-headed), while the bet F on the fair coin yields a one-half chance of the prize in each event.

Now observe that bet F is a 50:50 mixture of bets H and T because the chances of winning associated with F in each state of the world are an equal-weighted average of those associated with H and T. So A–A Independence requires that, if the agent is indifferent between bets H and T, then she should also be indifferent between bets H and F. This feature of her preferences, that they be unaffected by the spread of the chances, is not required by the Sure-Thing Principle, which imposes no constraints at all on how consequences should be valued. Nor can vN–M's Independence axiom be

TABLE 13.3. *Hedging Your Bets*

	Heads bias	Tails bias
<i>H</i>	1	0
<i>T</i>	0	1
<i>F</i>	0.5	0.5

applied here, since the chance-consequences appear in different states. So A–A Independence must impose constraints that exceed those jointly imposed by these other two conditions.

We are now finally in a position to characterise the kind of cautionary attitude we are interested in. A–A Independence requires indifference to the spread of chances; when an agent satisfies it, we will say that she is *ambiguity-neutral*. In contrast, if she prefers *g* to both *f* and *h* when indifferent between the latter two, we will say that she is *ambiguity-averse*. Formally, following Schmeidler (1989):

Ambiguity Aversion If $f \approx g$ then $\alpha f + (1 - \alpha)g \succsim f$

So ambiguity aversion on this account is a cautionary attitude canonically revealed in a preference for hedging against spreads of chances. My task now is to assess its significance for our understanding of rational decision making in conditions of severe uncertainty, looking first at competing models of ambiguity aversion and then at the question of whether it is rational.

13.5 MODELS OF AMBIGUITY AVERSION*

Ambiguity aversion plays a crucial role in the characterisation of the rules of cautious decision making presented in the previous section; most notably in the representation theorem of Itzhak Gilboa and David Schmeidler (1989) for the MMEU rule. Gilboa and Schmeidler adopt all the Anscombe and Aumann axioms, with the exception of A–A Independence, which they replace with a weaker version – called C-Independence – which restricts the separability requirement to mixtures with constant acts. More precisely, it requires that, for all acts $f, g \in \mathcal{F}$ and constant acts $\bar{h} \in \mathcal{H}$ and for all $\lambda \in [0, 1]$

C-Independence $\lambda f + (1 - \lambda)\bar{h} \succsim \lambda g + (1 - \lambda)\bar{h} \Leftrightarrow f \succsim g$

What they then prove is that an agent who satisfies these axioms plus ambiguity aversion can be represented as making choices in accordance with the MMEU rule relative to a set of probability functions, $C = \{P_1, \dots, P_n\}$, and a utility function *U* on consequences – i.e. they can be represented as agents with

imprecise degrees of belief C who maximise minimum expected utility relative to C in their choice of action.

Gilboa and Schmeidler's result has been used to support two different claims: an explanatory one and a normative one. The first is the claim that the hypothesis that agents maximise the minimum expected utility in their choices provides the best explanation of ambiguity-averse behaviour. The second is the claim that rationality requires agents with imprecise degrees of belief to maximise the minimum expected utility relative to these beliefs. An evaluation of these claims depends on two distinct issues. Firstly, whether the axioms adopted by Gilboa and Schmeidler are both necessary and sufficient for rationality (an issue of most relevance to the normative claim). And, secondly, whether the numerical functions that represent these preferences are indeed measures of the agent's degrees of belief and desire; in particular, whether C truly measures her imprecise degrees of belief. For their result establishes only that agents whose choices conform with their axioms are maximising minimum expected utility relative to *some* set of probabilities. Those with behaviourist inclinations will dismiss the second question on the grounds that the agent's 'true' beliefs are simply those that are revealed in her choices or, more radically, that her beliefs are nothing other than constructions out of her behaviour. But, even if we grant this, what grounds do we have for basing the construction on the MMEU rule? The fact that this construction works is not sufficient to establish the explanatory claim. It must be shown that there is no other way of constructing an agent's degrees of belief and desire from her preferences that affords an equally adequate explanation for them.

In fact, however, there are. The most salient one for our purposes is the smooth ambiguity model of Peter Klibanoff, Massimo Marinacci and Sujoy Mukerji 2005 (hereafter KMM), which provides a rival characterisation of ambiguity-averse choice. The smooth ambiguity model too is backed up with a representation theorem for acts with lottery consequences that shows that, if an agent's preferences over such acts satisfy Savage's axioms and her preferences over constant acts (those that can be identified with the constant lottery they determine in each state of the world) satisfy the vN–M axioms, then her preferences can be represented by a probability P defined on the extended Boolean algebra of prospects, and a utility defined on lottery consequences, such that for all acts f and g

$$f \succeq g \Leftrightarrow \sum_i \phi(\mathbb{E}_i(f)) \cdot P(S_i) \geq \sum_i \phi(\mathbb{E}_i(g)) \cdot P(S_i) \quad (13.1)$$

where $\phi : \Re \rightarrow \Re$ is a mapping on the real numbers representing the agent's attitudes to ambiguity and $\mathbb{E}_i(f)$ is the vN–M expected utility of f in state S_i . Crucially, if her preferences also satisfy ambiguity aversion then ϕ will be a concave transformation of the expected utilities.

On KMM's smooth ambiguity model the value of an act is an expectation based on (subjective) probabilities for (objective) probabilities – the chances of

the goods at stake – with the notable feature that the two tiers of probability are not reducible to a single expectation for these goods. This model generalises the one I proposed in Chapter 9, which applied only to decision problems (such as the Ellsberg Paradox) in which only one kind of good was at stake. In this case, the $vN-M$ expected utility of a consequence reduces to the desirability of the chance of this single good and so the chances themselves serve as a $vN-M$ index for preferences over the chances of such goods (because the greater the chances, the better). In the KMM model, the transformation is applied to the $vN-M$ expected utilities without reference to the types of good determining them. This seems to be reasonable when the goods are just different quantities of money, but not when the chances concern very different types of goods, such as health and money, with regard to which agents may reasonably have rather different degrees of ambiguity aversion. In these cases, the desirabilities of the chances of the different goods must be determined first and only then aggregated.

My main disagreement with KMM, however, concerns the interpretation of the parameter φ representing the agent's uncertainty attitudes to chances. KMM take $vN-M$ expected utility to be an appropriate measure of the desirability of chances (more generally, of lotteries) and ϕ as a transformation induced by the kind of epistemic attitude that Ellsberg postulates: a dislike of a lack of information that distorts the subjective probabilistic weighting of outcomes. I view $\phi(\mathbb{E}_i(\cdot))$ itself as the correct measure of the desirability of the lotteries, with ϕ a pragmatic attitude to uncertainty about chances that is encoded in the concavity of the utility function for chances. So interpreted, the smooth ambiguity model provides an explanation for ambiguity aversion that is perfectly consistent with Bayesian norms of rationality.

The upshot is that there is more than one model of ambiguity aversion on the table. I don't think we are in a position to definitively settle the question of which model is the most adequate descriptively and normatively, but I do think it's clear that MMEU is not it. Although it explains the ambiguity-averse pair of preferences $L_4 \succ L_3$ and $L_1 \succ L_2$ in the Ellsberg set-up (Table 13.2), it also implies that agents facing this decision problem will be indifferent between L_3 and L_1 . Since L_3 weakly dominates L_1 , this implication is both descriptively implausible and normatively unsatisfactory. The α -MEU rule does not have this implication but, like all rules that take into account only the maximum and minimum expected utilities of each action, does not discriminate between actions with quite different intermediate possibilities. We can illustrate this using the Ellsberg set-up once again, but dropping the assumption that the proportion of red balls is given, so that the situation is one of complete ignorance about the chances of drawing a ball of any particular colour. In this case, L_2 and L_4 have the same maximum and minimum expected utilities (respectively, the utility of \$0 and of \$100) and so must be regarded indifferently under the α -MEU rule. L_4 weakly dominates L_2 , however, and I doubt anyone would choose L_2 in these circumstances. Nor should they.

In contrast, the smooth ambiguity model does, it seems to me, offer a plausible account of rational choice in the kinds of contexts exemplified by the Ellsberg Paradox, in which the state space is given in the description of the problem and in which the symmetries in the problem make application of the Principle of Indifference natural. In other contexts, those previously labelled as Grade 3 uncertainty in contrast to Grade 2 ambiguity, in which such symmetries are absent, the smooth ambiguity model is less compelling. For in such cases we have little to guide us in assigning subjective probabilities to the chances and no reason to trust them. In the next chapter I will defend a different model for such cases.

13.6 THE RATIONALITY OF AMBIGUITY AVERSION

Ambiguity aversion leads to decision making that is relatively cautious in the sense that ambiguity-averse agents will prefer actions with narrower spreads in the chances of outcomes. There is little doubt that agents do display caution of this kind, but is it rational? A number of philosophers and economists have argued that it is not (see, for instance, Adam Elga, 2010, and Nabil Al-Najjar and Jonathan Weinstein, 2009), others that it is (see, for instance, Levi, 1990, and Gilboa, 2009). In this section I will defend the latter position, but also stress the costs ambiguity aversion can impose on the agent.

Let's start with risk aversion by way of comparison. An agent is canonically said to be risk-neutral with respect to some divisible good if she is indifferent between a fixed amount of the good and a lottery which yields the same expected amount of it, but risk-averse (-loving) if she prefers the former (latter). For instance, someone who is risk-averse with respect to money will prefer \$50 to a lottery with a 50:50 chance of paying out either \$100 or nothing. More generally, risk aversion with respect to a divisible good G is manifested in a preference for a fixed quantity g of the good to a lottery with an expected return of g .

In a similar fashion, we can define uncertainty attitudes to goods. Consider Table 13.4, displaying prospects $L1$, $L2$ and $L3$, which have monetary consequences that depend on the truth or falsity of some event E . Suppose that an agent is indifferent between $L2$ and $L3$, thereby revealing that she regards E as likely to be true as not (assuming that the monetary consequences have desirabilities that are independent of E). Then we can say that she is neutral to the uncertainty regarding the monetary consequences if she is also indifferent between $L1$ and $L2$ and between $L1$ and $L3$; uncertainty-averse regarding money if she prefers $L1$ to both $L2$ and $L3$; and uncertainty-loving regarding money if she prefers both $L2$ and $L3$ to $L1$. These attitudes reflect, in essence, different desirability functions for quantities of money.

Now consider someone who is uncertainty-averse with respect to the *chances* of receiving some good (divisible or otherwise). She will prefer acts which yield constant chances of getting the good over acts with the same

TABLE 13.4. *Uncertainty
Aversion regarding Money*

	<i>E</i>	$\neg E$
<i>L1</i>	\$50	\$50
<i>L2</i>	0	\$100
<i>L3</i>	\$100	0

TABLE 13.5. *Hedging Chances*

	<i>RBB</i>	<i>RBY</i>	<i>RYY</i>
<i>B</i> ₀	$\frac{2}{3}$	$\frac{1}{3}$	0
<i>B</i> ₁	$\frac{1}{3}$	$\frac{1}{3}$	$\frac{1}{3}$
<i>B</i> ₂	0	$\frac{1}{3}$	$\frac{2}{3}$
—	—	—	—
<i>B</i> ₃	$\frac{1}{3}$	$\frac{2}{3}$	1
<i>B</i> ₄	$\frac{2}{3}$	$\frac{2}{3}$	$\frac{2}{3}$
<i>B</i> ₅	1	$\frac{2}{3}$	$\frac{1}{3}$

expected chances when the chances vary by state of the world. Consider Table 13.5, for instance, which adds two acts (*B*₀ and *B*₅) to the simplified Ellsberg set-up presented before. Someone who regards the distributions *RBB* and *RYY* as equiprobable will be indifferent between *B*₂ and *B*₀ and between *B*₃ and *B*₅. If, furthermore, she is neutral with regard to the chances of monetary gain that are the outcomes of these acts, she will regard *B*₁ as equally good as both *B*₂ and *B*₀ and *B*₄ as equally good as both *B*₃ and *B*₅. But if she is uncertainty-averse with respect to these chances she will prefer *B*₁ over the other two because the desirability difference between a chance of one-third of the \$100 and no chance of it exceeds that between a chance of two-thirds and a chance of one-third. Similarly, she will prefer *B*₄ to both *B*₃ and *B*₅.

Uncertainty attitudes to goods and uncertainty attitudes to the chances of these goods are logically independent. One could be uncertainty-neutral with regard to money, but uncertainty-averse with respect to the chances of obtaining it. Or just the other way around. But in one crucial respect they are similar: on the face of it, there is nothing particularly rational or irrational about having one uncertainty attitude rather than another. We certainly do, as a matter of fact, care about the chances of outcomes as well as the outcomes themselves. There is a difference, we tend to think, between having no lottery ticket at all and having a lottery ticket which is not, in fact, a winner. And between succeeding at a task when the chance of doing so was low and

TABLE 13.6. *Cost of Hedging*

Options	Composition of urn A				
	H_0	H_1	H_2	...	H_{100}
A	0	0.01	0.02	...	1
B	1	0.99	0.98	...	0
C	0.5	0.5	0.5	...	0.5

succeeding at it when the chance of doing so was very high. (This is just what the Chapter 9 examples of Ann the mountaineer and Bob the surly sibling revealed.)

Now, on my account, ambiguity aversion, as characterised by a preference for hedging, can be rationalised by aversion to uncertainty about chances. So the caution that it induces is perfectly rational from the perspective of the moderate Humeanism informing this book (which passes no judgement on the content of desires, only on their consistency). On the other hand, there *is* a price to be paid for such an aversion to spreads of chances. For an agent who prefers to hedge her chances will be willing to pay for such an opportunity even when it does not improve her expected gains. To see this, consider a close variant of the example that was presented as a challenge to Imprecise Bayesianism in Chapter 11.

Suppose 200 balls, 100 white and 100 black, are divided between two urns, respectively labelled A and B, with any ratio of the colours in the urns. You will be given an opportunity to bet on a white ball being drawn; the bet costs \$10 and you win \$50 if the drawn ball is white. There are three options you must choose amongst: A, a bet a white ball drawn from urn A; B, a bet on on a white ball drawn from urn B; and C, a bet on a white ball drawn from either urn A or urn B depending on the result of a toss of a fair coin. These options are displayed in Table 13.6, with the cell entries recording the chances of winning a prize for each hypothesis as to the composition of urn A for each of the options.

Table 13.6 makes evident the fact that option C offers the opportunity to hedge on the chances yielded by options A and B. So an ambiguity-averse agent should be willing to pay some amount of money to secure it in preference to options A and B. Indeed, if she is very averse to ambiguity she will be willing to buy bet C but not A or B. She is now open to exploitation, however, for someone could sell her bet C, toss a coin and then, once it has landed, reveal which urn the ball will be drawn from. At this point the agent in effect holds either bet A or bet B. Should she not be willing to sell these bets back for some further amount, leaving her out of pocket and back where she started?

Not necessarily. The fact that she would not buy either bet at a particular price does not mean that she is willing to sell them at that price. For her

position as buyer or seller is rather different. As a buyer she trades a fixed amount of money for the chance of a gain. As a seller she acquires a fixed amount for the chance of a loss. If her attitudes to chances of gains and losses are different then she will not view sales and purchases symmetrically.¹¹ And so she might refuse to sell the bet back. Nonetheless, it does not seem irrational to view chances of gains and losses in the same way. In which case, an ambiguity-averse agent will be vulnerable to exploitation of this kind.

Does this make ambiguity aversion irrational? I don't think so. The way to think about it is in analogy with someone whose preferences change in a predictable way. Suppose that on Monday I find the prospect of going to the opera on Saturday sufficiently attractive that I am willing to purchase a ticket at £200. By Friday, however, I am feeling tired and the prospect of a late night seems unappealing. I now wish that I had not bought the ticket and would be willing to sell it back for less than £200. This fact about me makes me vulnerable, for someone who could anticipate my attitude changes could 'exploit' me by first selling me the ticket and then buying it back, leaving me out of pocket. (Of course, if I predict that I will end up selling the ticket back, I should not purchase it in the first place. So too the ambiguity-averse agent should not purchase bet C if she knows that the coin toss will be revealed to her, since she can anticipate the effect of the information about how the coin has landed on her attitudes. Sophisticated agents will not allow themselves to be exploited.)

There is no doubt that to have preferences that can be exploited in this way can be detrimental. If I could change them, that would be to my advantage, but, in reality, transforming preferences can be expensive – often prohibitively so. In any case, instability of preference is not irrationality; nor, more generally, is vulnerability to exploitation a sure sign of it. So the possibility of exploitation does not in itself show that ambiguity aversion is irrational. The ambiguity-averse agent's preferences imposes costs on her that an ambiguity-neutral agent does not face. But, like agents with expensive tastes, ambiguity-averse agents simply have to do the best they can with the preferences they are endowed with.

13.7 CONCLUDING REMARKS

I set out in this chapter to assess how an Imprecise Bayesian might make decisions in situations of severe uncertainty in which she lacks precise probability and/or desirability values for some of the prospects relevant to the evaluation of her options. Two strategies were proposed: furnishing the required probabilities and desirabilities by reaching a subjective judgement on the basis of whatever information she holds plus various non-evidential

¹¹ There is some evidence that people do view chances of gains and losses differently. See Wakker (2010) for the evidence and Bradley (2016) for a discussion.

considerations; and adopting a decision rule that does not require precise probabilities and desirabilities as inputs. The question that now needs to be addressed more explicitly is whether pursuit of either of these strategies risks confrontation with any of the core Bayesian rationality principles and, if so, what lesson is to be drawn from this fact. The question is a rather complex one, and evidently there are not only many variants of the two strategies to consider but also many nuances to the issue of what it would mean to abandon rather than modify the Bayesian approach. But there are nonetheless some preliminary conclusions to be drawn from the discussion.

Pursuit of the ‘making up one’s mind’ strategy is of course required by Classical Bayesianism in order that expected utility maximisation be applicable. But it is also compatible with the kind of subjectivist Imprecise Bayesianism that I argued for in the previous chapter, which permits precision beyond what is required by the evidence. The question remains, however, whether there are principled ways of doing so that are consistent with Bayesian principles. The outlook in this regard is quite mixed. On the positive side, although I don’t think much of the claim that MaxEnt is the uniquely rational way of forming precise probabilities, it does seem to me to offer the subjectivist a useful tool for arriving at a judgement in the kind of circumstance that I called Grade 2 uncertainty, when the state space and its symmetries are given by the description of the problem. (The Ellsberg Paradox is a prime example of such a circumstance; hence the importance of explaining how an agent might employ the Principle of Indifference to assign probabilities to relevant chance hypotheses and still act cautiously.) The flip side of this is that in situations of Grade 3 uncertainty, when the state space is not given, the grounds for reasonable application of MaxEnt are lacking, as it is no longer possible to identify symmetries in a non-arbitrary way. In such circumstances we might look to other considerations, such as confidence, to evaluate possible judgements. But the most commonly proposed version of this, linear averaging by the application of confidence weights on probabilities, turned out to face a number of grave difficulties. Whatever its pragmatic merits it does not provide a method for reaching a judgement fully consistent with Bayesian norms. Nonetheless, there is something right about this approach, and, in the next chapter, I will offer an alternative method for applying considerations of confidence to judgement.

In contrast to the first strategy, the second has typically been interpreted as one of providing *rival* decision rules to the Bayesian ones and hence as calling into question the descriptive and/or normative validity of Bayesian decision theory. I see things somewhat differently. There is an important difference between proposing additional considerations to those characteristics of Bayesianism and applying them in ways which require violation of the Bayesian norms of rationality. Since completeness is not a requirement of rationality, the core Bayesian rationality conditions on preference, together with the broad requirement of preference-based choice, do not completely

determine what choices an agent should make. So a Bayesian must apply additional considerations – such as caution, confidence, flexibility and robustness – to bridge the gap between preference and choice. But doing so should not require violations of the rationality requirements on incomplete preferences.

Tension between the two arises only when such considerations are allowed to trump the Bayesian norms. For example, if caution is applied in the manner encoded in the MMEU decision rule, then the preferences over actions it supports may violate the Sure-Thing Principle. On the whole, however, I think this counts against applying caution in this way. Nor does the permissibility of ambiguity aversion require that we do so, as caution of this kind is perfectly consistent with a broad-minded Bayesianism that countenances a variety of attitudes to chances. It is true that ambiguity aversion is not consistent with the combination of Bayesianism and the treatment of risk given by the vN–M theory, but I have already argued that the latter is too restrictive. So this does not provide compelling grounds for allowing violations of Bayesian norms for incomplete preferences.

All this provides no more than a partial answer to the question of how an Imprecise Bayesian should make decisions. For the kinds of situations exemplified by the Ellsberg Paradox display only Grade 2 uncertainty, and the fact that we are able to provide a coherent decision theory for them does not mean that we have a proper handle on situations of Grade 3 or severe uncertainty. Our task in the final chapter will be to address this problem.