

Belief and Credence

WENG HONG TANG

A THESIS SUBMITTED
FOR THE DEGREE OF DOCTOR OF PHILOSOPHY
OF THE AUSTRALIAN NATIONAL UNIVERSITY
DECEMBER 2009

Declaration

This dissertation is solely the work of its author. No part of it has been submitted for any degree or is currently being submitted for any other degree. To the best of my knowledge, any help received in preparing this dissertation, and all sources used, have been acknowledged.

Acknowledgements

Thanks to Ole Koksvik, Jo-Anne Weinman, and Tug for shelter and good company during the last six weeks of my dissertation writing.

Thanks to the following people for discussion and comments, provided either one-on-one or at talks: Ben Blumson, David Bourget, David Chalmers, Mark D’Cruz, Kenny Easwaran, Lina Eriksson, Frank Jackson, Brian Kim, Leon Leontyev, Stephan Leuenberger, Aidan Lyon, John Matthewson, Leone Miller, Jonathan Schaffer, and David Wall.

Many thanks to John Cusbert, Clement Huang, Ole Koksvik, Holly Lawford-Smith, and Declan Smithies for suffering through drafts of various chapters and for very useful comments.

Finally, I’m indebted to Jens Christian Bjerring, Alan Hájek, and Wolfgang Schwarz for invaluable discussion and for extensive feedback on the dissertation from start to finish.

Abstract

There are two main tasks I aim to accomplish in this dissertation: first, show that credences rather than (binary) beliefs do all the heavy lifting in the explanation and prediction of action and behaviour, and second, show how the notions of credence and belief are related. In chapter 1, I go through some preliminaries. In chapter 2, I consider some views according to which beliefs about objective probabilities can do all the work that credences do. I argue that such views are false. In chapter 3, I put forward certain constraints on rational credences that are more realistic than those provided by the probability axioms. Rational credences are typically taken to be subjective probabilities, but according to some, this requires that agents be logically omniscient. Thus, one may worry that we cannot appeal to credences to explain and predict the action and behaviour of rational but non-ideal agents like human beings. The more realistic constraints I put forward are meant to assuage such a worry. In chapter 4, I consider some arguments for thinking that beliefs have some role to play in explaining and predicting action and behaviour that credences cannot play. Some of these arguments aim to show that credences are conceptually dependent upon beliefs, whereas some aim to show that beliefs have a non-redundant role to play in explaining our acts of reasoning and assertion. I argue that these arguments are unsound. However, even if credences do all the heavy lifting in explaining and predicting action and behaviour, one might think that beliefs still have their pride of place in *epistemology*. For one thing, knowledge is thought to entail rational *belief*. As we shall see, rational beliefs

also have a certain role to play with respect to the *epistemic* evaluation of reasoning and assertion. But if belief can be cashed out in terms of credence (perhaps plus some other notions such as desire), then it seems that credences can take the place of beliefs even in epistemology. Thus it will be worth investigating how the notions of belief and credence are related. In chapter 5, I consider a dilemma that confronts attempts to relate the two notions. Presumably, to believe or for one to be rational in believing a proposition, one has to have a high enough credence in it. If we hold that the minimum credence needed for belief or rational belief is 1, then we will have rather few beliefs or rational beliefs. But if we hold that the threshold is less than 1, then we face a host of other worries—among other things, we will have to contend with the lottery paradox, and with the failure of certain closure principles. In chapter 6, I argue that extant accounts of the relationship between belief and credence are unable to resolve the dilemma. In chapter 7, I defend my account of how the two notions are related, and argue that it resolves the dilemma. I conclude the dissertation in chapter 8.

Contents

1	Preliminaries	1
1.1	Credences: A Philosopher of Mind’s Perspective	2
1.2	Credences: A Decision Theorist’s Perspective	6
1.3	Plan	9
2	No Doing Without Credences	12
2.1	Doing Without Credences?	12
2.2	Problems with Subjective Probabilities?	16
2.3	All-or-Nothing Beliefs about Objective Probabilities	18
2.3.1	The Replacement Thesis	18
2.3.2	Different Kinds of Probability	19
2.3.3	Christensen’s Argument against the Replacement Thesis	21
2.3.4	Other Problems for the Replacement Thesis	24
3	Credences and Logic	31
3.1	The Probability Axioms	31
3.2	Problems with Probabilism	36
3.2.1	Too Demanding	37
3.2.2	Too Beholden to Logic	42
3.2.3	Does Scant Justice to Coherence	43
3.2.4	Does Not Provide a Good Account of Action and Behaviour	44
3.2.5	Does not Respect the Implication/ Inference Distinction	45

3.3	Proposed Constraints on Rational Credences	46
3.4	Objections and Replies	54
3.4.1	Irrationality Twice Over	55
3.4.2	Too Undemanding	55
3.4.3	Lack of Mathematical Structure	57
4	Credences Do All the Heavy Lifting	63
4.1	The Sense in which Credences Do the Heavy Lifting	63
4.2	The Supposed Dependence of Credence on Binary Beliefs	67
4.2.1	Levi's view	67
4.2.2	Nozick's view	70
4.2.3	Lance's view	74
4.3	Binary Beliefs and their Putative Role in Reasoning	78
4.3.1	Practical reasoning	78
4.3.2	Theoretical reasoning	88
4.4	Binary Beliefs and Assertion	89
4.5	Binary Beliefs and Epistemology	90
5	Belief and Credence: A Dilemma	93
5.1	The Epistemic Notion of Binary Belief	93
5.2	The Threshold View and the Lockean Thesis	96
5.3	The First Horn of the Dilemma: Belief and Full Credence	98
5.4	The Second Horn of the Dilemma	99
5.4.1	Arbitrary Thresholds	100
5.4.2	Descartes's injunction	100
5.4.3	The Lottery Paradox	102
5.4.4	Closure	105
5.4.5	Belief and confidence	119

6	Resolving the Dilemma I	123
6.1	Seizing the First Horn of the Dilemma	123
6.2	The Contextualised Threshold View	129
6.3	The Contextualised Lockean Thesis	131
6.4	Proposition-relative Thresholds	133
6.5	Beliefs, Acts, and Preferences	136
6.5.1	Kyburg's version of Belief-Act	138
6.5.2	Fantl and McGrath's version of Belief-Act	140
6.5.3	Weatherson's version of Belief-Preference	141
6.5.4	Fantl and McGrath's version of Belief-Preference	142
6.5.5	Some Problems	143
7	Resolving the Dilemma II	151
7.1	Two Theses	151
7.1.1	Whose Context?	155
7.1.2	Presupposition	158
7.1.3	Some Implications	161
7.2	Applying the Theses	162
7.2.1	Dealing with Descartes's Injunction	163
7.2.2	Dealing with Lotteries	163
7.2.3	Dealing with Closure	164
7.2.4	Dealing with Confidence Claims	167
7.3	A Comparison with Similar Views	168
7.4	Objections and Replies	171
7.4.1	The Threat of Conditional Bets	171
7.4.2	Still Too Few Beliefs and Rational Beliefs	175
7.4.3	Moving the Bump Under the Rug?	177
7.4.4	Animals and Beliefs	178

7.4.5	Retraction Worries and Mixed Situations	180
7.4.6	Beliefs and Context-Relativity	183
8	Conclusion	188
	Bibliography	192

Chapter 1

Preliminaries

We often appeal to our beliefs and desires to explain or predict the way we act and behave. In such explanation or prediction, it is natural to think that beliefs and desires come in various strengths. Why did Ida put on a scarf before leaving the house? Well, one explanation is that she desires not to catch a cold, and she believes both that it is chilly outside, and that wearing a scarf when it is chilly outside will decrease her chances of catching a cold. But suppose someone offers Ida a thousand dollars to go outside without a scarf, and she accepts the offer. How do we explain this? If desires were just all-or-nothing, we would be hard pressed to explain Ida's acceptance of the offer. However, desires vary in strength—Ida desires a thousand dollars more strongly than she desires to avoid catching a cold. Now suppose someone offers Ida the option to bet either on it being chilly outside or on the sun rising tomorrow, and she chooses to bet on the latter. How do we explain this? If beliefs were just all-or-nothing, we would be hard pressed to explain Ida's choice. However, beliefs vary in strength—Ida believes that the sun will rise tomorrow more strongly than she believes that it is chilly outside.

Beliefs that are all-or-nothing are also known as binary beliefs, flat out beliefs, black-or-white beliefs, outright beliefs, or full beliefs.¹ Beliefs that

¹'Full belief' is sometimes used to refer to a belief of maximum strength. For this reason, I shall avoid referring to binary beliefs as full beliefs, unless the authors I am discussing do so.

come in varying strengths are also known as credences, degrees of belief, graded beliefs, degrees of confidence, or partial beliefs. I shall sometimes use ‘belief’ to mean ‘binary belief’, and sometimes use it in a way that is neutral between the binary sense and the graded sense of the word. The context will make it clear how I am using it.

What is the relationship between credence and binary belief? This dissertation is an attempt to answer the preceding question. There are two main tasks I aim to accomplish. First, I aim to show that credences rather than beliefs do all the heavy lifting in the explanation and prediction of action and behaviour. Second, I aim to show how the notion of (rational) credence and the notion of (rational) binary belief are related. But before all that, let me say more about beliefs (in particular credences), first from the perspective of the philosophy of mind, and then from the perspective of decision theory.

1.1 Credences: A Philosopher of Mind’s Perspective

The folk psychological picture of beliefs and desires takes them to be intimately related to action and behaviour. If one believes something, then one is disposed to act and behave in certain ways, or to enter into certain mental states, given one’s desires and other mental states. By behaviour, I mean what David Lewis (1974) means by behaviour, namely, ‘raw behaviour—body movements and the like’ (p. 338). By action, I include ‘behaviour specified partly in terms of the agent’s intentions’ (*ibid.*).² But I also take it to include mental actions, such as the entertaining of a thought, the inner assertion of some sentence, or the painting of a mental image. For example, if one is feeling sad and desires to be cheered up, and one believes that the mental image of a pink elephant doing a pirouette will cheer one up, one might be disposed to

²When quoting an author, symbols or the spelling of certain words may be modified (silently) to preserve stylistic consistency. I shall also ignore the use-mention distinction if there is no risk of confusion.

paint such an image in one's mind.

I take the folk psychological explanation and prediction of action and behaviour to be reliable by and large, and I take such reliability to be evidence that beliefs and desires are real, or at least, that it is useful to talk as if they are. As such, I am concerned with an issue that is internal to those who are *realists* or at least *instrumentalists* about such psychological entities. Given that beliefs are real, or at least, that it is useful talk as if they are, do credence and binary belief both have their place in the explanation and prediction of action and behaviour, or does one threaten to displace the other?

Various views about belief are consistent with the folk psychological picture of belief sketched earlier. Suppose that Ida, who believes that it is cold outside, puts on a scarf before leaving the house. According to a *functionalist* such as Lewis (1994), in believing that it is cold outside, Ida is in a robust mental state that, in response to sensory stimuli and other mental states such as her desires, typically disposes her to put on a scarf. According to a *dispositionalist* such as Eric Schwitzgebel (2002), for Ida to believe that it is cold outside is 'nothing more or less than being disposed to do and experience certain kinds of things', e.g. to put on a scarf (p. 250). According to an *interpretationist* such as Daniel Dennett (1991), for Ida to believe that it is cold outside is just for us to adopt what he calls the *intentional stance* and attribute to Ida such a belief in order to predict her behaviour. Both the dispositionalist and the interpretationist may also be instrumentalists. They may hold that although beliefs are not robust mental states and are not real in the way that a rock is real, it is nonetheless useful for us to appeal to them in explaining and predicting the way agents such as human beings behave, just as it is useful for us to appeal to centres of gravity in predicting whether vases placed at the edge of a table will topple over. In discussing the relationship between binary belief and credence, I shall remain neutral as to the truth of

each view mentioned above.³ Of course, if one is an *eliminativist* about beliefs like Paul Churhland (1981), and holds not only that there are no beliefs, but that we ought not to engage in belief talk, then the question of how binary belief is related to credence will be of scant interest. This dissertation speaks to those who do not subscribe to such a position.

The folk psychological picture is also consistent with certain distinctions between different kinds of belief. Some philosophers distinguish between *explicit* and *implicit* beliefs. An explicit belief is a belief that is explicitly represented in the mind, perhaps by being stored somewhere in the brain. An implicit belief is one that can be easily inferred from one's explicit beliefs, and is not stored in the brain the way that explicit beliefs are—there are just too many implicit beliefs. Suppose that Ida has an explicit belief that she has published no papers. Based on this explicit belief, she might have the implicit belief that she has published no more than one paper, and the implicit belief that she has published no more than two papers, and so on.

It is debatable whether there are really any beliefs that are explicitly represented in the mind, or stored in the brain.⁴ But suppose that there are, and that there exists a real distinction between explicit and implicit beliefs. When I claim that credences rather than binary beliefs do all the heavy lifting, I intend the claim to be about credences and binary beliefs, both explicit and implicit. For if there are both explicit and implicit beliefs, we will appeal to both in the explanation and prediction of action and behaviour. Suppose Ida reads an advertisement for a job that says that those who have published no papers need not apply. Presumably, Ida will not apply. Why? Well, one explanation is that she has the explicit belief that she has published no papers. And presumably, if she had read a different advertisement saying that those

³For a discussion of such views, see Schwitzgebel (2008).

⁴Fodor (1989) and Harman (1970) endorse such a view of belief, whereas Dennett (1978) denies it (p. 22; p. 34; pp. 104-107).

who have published no more than two papers need not apply, she would not have applied. For she believes implicitly, even if not explicitly, that she has published no more than two papers.

One may also distinguish between *occurrent* beliefs and *non-occurrent* beliefs. An *occurrent* belief is one that is somehow before one's mind, whereas a *non-occurrent* belief is one that is not. For example, one's belief that Canberra is the capital of Australia is *occurrent* when one holds the thought that Canberra is the capital of Australia in mind. One's belief that European magpies are of a different species from Australian magpies is *non-occurrent* when one has the belief while sleeping. We might think of an *occurrent* belief as one that is ready to be deployed in conscious deliberation or reasoning, whereas a *non-occurrent* belief has to be brought to mind before being so deployed. The distinction between *occurrent* and *non-occurrent* beliefs is not the same as that between *explicit* and *implicit* beliefs. *Explicit* beliefs can be either *occurrent* or *non-occurrent*—Ida's *explicit* belief that she has published nothing can either be present before her mind or not. Even when she is sleeping, this belief of hers still counts as *explicit* if it is stored somewhere in her brain. *Implicit* beliefs, however, can only be *non-occurrent* if *occurrent* beliefs have to be explicitly represented in the mind to be readily deployable in conscious deliberation and reasoning.

I have stressed that our beliefs and desires are intimately connected with the way we act. The claim is a descriptive one. But as several philosophers have pointed out, this is a case in which the descriptive is tinted with the normative—for us to ascribe to an agent any beliefs or desires at all, we must impute to the agent some amount of rationality.⁵ If we have such-and-such beliefs and such-and-such desires, then we are disposed to act or behave in

⁵For example, see Cherniak (1981), Davidson (1985), Dennett (1978), Harman (1986) and Lewis (1974) (p. 161; pp. 89-90; pp. 9-10; p. 7; p. 337).

a certain way. What kind of way? Presumably, one that serves our desires according to our beliefs, or at least, serves them well enough. But to act or behave in such a way is to possess at least a modicum of rationality. In *describing* how Ida's beliefs and desires combine to yield putting-on-a-scarf behaviour, we also *rationalise* it.⁶ It will be strange, for instance, if Ida desires not to catch a cold, believes that she is likely to catch a cold if she wanders outside without wearing a scarf, and yet does precisely that in the absence of countervailing factors (like the lure of a thousand dollars for not wearing a scarf). We will find such behaviour inexplicable, because we assume that Ida is more or less rational, and will not act in a way that is incongruous with her beliefs and desires.

Granted, we do not always act or behave in ways that serve our desires according to our beliefs—we are cognitively bounded, and we sometimes suffer from lapses of rationality, get distracted, or undergo memory loss. But we act and behave more or less rationally more often than not, or it would be hard for us to make sense of one another. At the very least, if we were to be allocated more cognitive resources, be fully attentive, and enjoy perfect memory, then we would be disposed to act in ways that serve our desires well enough.

1.2 Credences: A Decision Theorist's Perspective

Talk of acting in ways that serve our desires according to our beliefs can be made more precise by turning to subjective decision theory. The theory provides us with a certain set of formal tools to help us model beliefs and desires, or to help us prescribe norms for action and behaviour. Let $P(\cdot)$ be Jill's subjective probability function, representing her credences. Let $U(\cdot)$ be her utility function, representing her degrees of desire. Let the contents of credences and desires be propositions, whatever propositions might turn out

⁶This way of putting the point is inspired by Davidson (1985)—see pp. 89-90.

to be. And let the various propositions that might obtain conditional on the performance of a particular act be represented by s_1, s_2, s_3, \dots . The expected utility of performing the act is given by $\sum_i P(s_i|a)U(s_i \& a)$, where a is the proposition that one performs the act in question. To perform an act that maximises expected utility is to perform an act that has an expected utility no smaller than that of any other act.⁷

Now subjective decision theory, construed as a *descriptive* enterprise, models how we act according to our beliefs and desires. But as mentioned earlier, the distinction between the descriptive and the normative is not as clear-cut as one might think, for the very concepts of belief and desire are inextricably tied to the concept of rational action. So some might see subjective decision theory as a model of *constitutive* rationality.⁸ On this view, to count as having beliefs and desires *at all* is to be disposed to act in ways that maximise (or at least satisfice) expected utility. Yet others have taken subjective decision theory to be a *prescriptive* enterprise that tells us that we ought to act in a way that maximises expected utility, whether or not we in fact act in such a way. Finally, one may think of decision theory as a mathematical toolkit that we can use to help us make decisions, in the way that one might use truth tables to check whether a certain line of reasoning is valid—call such a version of decision theory *toolkit decision theory*. The distinctions are worth keeping in mind, if only because some of the arguments that work against one version of decision theory need not work against other versions.⁹ In particular, as we shall see in chapter 4, some of the charges levelled against a descriptive version of decision theory are really only effective against toolkit decision theory.

For my purposes, I shall mainly be invoking the descriptive and pre-

⁷For simplicity, I shall work with evidential decision theory rather than with causal decision theory. Nothing much hangs on this choice.

⁸See, for example, Lewis (1986), p. 36.

⁹One may subscribe to various combinations of the different versions of subjective decision theory I have just sketched, although some combinations are more natural than others.

descriptive versions of subjective decision theory (while granting that the descriptive may be tinted with the normative). But one may worry that the theory would only be apt for describing or prescribing how severely idealised agents act and behave, or ought to act and behave. Subjective decision theory takes credences or rational credences to be subjective probabilities. But, as we shall see in chapter 3, some think that this leads to the result that agents or rational agents are logically omniscient. For instance, one of the axioms of probability theory says that for any logical truth t , the probability of t is 1. But then on the face of it, only the credences of idealised agents who are logically omniscient will satisfy such an axiom.

We can, however, preserve the spirit if not the letter of decision theory by looking for more realistic constraints on credences than those provided by the probability axioms. (I explore this in chapter 3.) Doing so promises to yield an unorthodox version of decision theory that is descriptively more accurate and prescriptively less demanding than traditional decision theory. But depending on our purposes, it may sometimes be expedient to invoke agents who are logically omniscient. In such cases, it will be proper to invoke subjective decision theory, as traditionally understood, to explain, predict, or evaluate the action and behaviour of such agents. (We shall encounter such a case in chapter 5.)

There are other issues concerning the use of probabilities to model credences. First, there is the issue of whether credences should always be represented by precise numbers.¹⁰ One might think that they should not. For example, one might think that credences can be vague. Or one might think that they are interval-valued or that they should be represented qualitatively or comparatively.¹¹ However, for the sake of simplicity and familiarity, I shall

¹⁰See Hájek and Smithson (forthcoming).

¹¹See, for example, Levi (1985), p. 390, and Fishburn (1981), pp. 144-145.

make the idealising assumption that credences always take on precise values. This would be an innocuous assumption for most of my purposes, but I will drop the assumption, and be explicit about doing so, in those cases in which the assumption tilts the case in favour of what I am arguing for.

Second, there are various unresolved paradoxes in decision theory such as the St. Petersburg Paradox and the Pasadena Paradox.¹² The Pasadena paradox, for example, purports to show that in certain cases in which an agent is faced with a decision problem, decision theory is silent on what the agent will do, or ought to do, even when it seems that it should not be silent. One cheerful response is to note that the troubling notion of infinity is somehow involved in the generation of such paradoxes, and then hold that since we do not have a good grip on infinity anyway, we should not worry about such paradoxes for the time being. Perhaps when we understand infinity better, we will see our way out of such (apparent) paradoxes. A less optimistic response is to acknowledge the problems, and admit that decision theory is imperfect. But until someone comes up with a better theory, we should continue using decision theory anyway, since it is the best thing we have. From a pragmatic point of view, there might not be any difference in adopting either of the two approaches. Since I think that at least one of the two approaches is the correct one to adopt, I shall set the paradoxes aside, and continue appealing to decision theory.

1.3 Plan

As we saw at the beginning of this chapter, there is, at least on the face of it, good reason to appeal to credences to explain and predict action and behaviour. Richard Jeffrey (1970) writes:

¹²Jeffrey (1983) discusses the St. Petersburg Paradox (pp. 151-155). The Pasadena Paradox was introduced by Nover and Hájek (2004).

our ordinary notion of [binary] *belief* is only vestigially present in the notion of degree of belief. I am inclined to think Ramsey [who made significant contributions to decision theory] sucked the marrow out of the ordinary notion, and used it to nourish a more adequate view. (p. 172; Jeffrey's emphasis)

Is Jeffrey right? There are some who think that he isn't, and that in fact, the notion of binary belief remains very healthy. I'm inclined towards Jeffrey's view, and shall show why in the rest of this dissertation.

Here's the plan. In chapter 2, I consider some views according to which binary beliefs about objective probabilities can do all the work that credences do. I argue that such views are false. In chapter 3, I put forward certain constraints on rational credences that are more realistic than those provided by the probability axioms. Rational credences are typically taken to be subjective probabilities, but according to some, this requires that agents be logically omniscient. Thus, one may worry that we cannot appeal to credences to explain and predict the action and behaviour of rational but non-ideal agents like human beings. The more realistic constraints I put forward are meant to assuage such a worry. In chapter 4, I consider some arguments for thinking that binary beliefs have some role to play that credences cannot play. Some of these arguments aim to show that credences are conceptually dependent upon binary beliefs, whereas some aim to show that binary beliefs have a non-redundant role to play in explaining our acts of reasoning and assertion. I argue that these arguments are unsound.

However, even if credences do all the heavy lifting in the explanation and prediction of action and behaviour, one might think that binary beliefs still have their pride of place in *epistemology*. For one thing, knowledge is thought to entail rational *belief*. As we shall see, rational beliefs also have a certain role to play with respect to the *epistemic* evaluation of reasoning and assertion.

But if belief can be cashed out in terms of credence (perhaps plus some other notions such as desire), then it seems that credences can take the place of beliefs even in epistemology. Thus it will be worth investigating how beliefs are related to credences. In chapter 5, I consider a dilemma that confronts attempts to relate the two notions. Presumably, to believe or for one to be rational in believing a proposition, one has to have a high enough credence in it. If we hold that the minimum threshold for belief or rational belief is a credence of 1, then we will have rather few beliefs or rational beliefs. But if we hold that the threshold is less than 1, then we face a host of other worries—among other things, we will have to contend with the lottery paradox, and with the failure of certain closure principles. In chapter 6, I argue that extant accounts of the relationship between binary belief and credence are unable to resolve the dilemma. In chapter 7, I defend my account of how the two notions are related, and argue that it resolves the dilemma. I conclude the dissertation in chapter 8.

Chapter 2

No Doing Without Credences

2.1 Doing Without Credences?

Can binary beliefs do the work that credences are supposed to do in the explanation and prediction of action and behaviour? Some philosophers think so. In particular, they hold that whenever it seems that we need to appeal to credences, we can instead appeal to binary beliefs about *objective probabilities*.¹ In this chapter, I shall examine and argue against such views.

Suppose that Jack has to decide which path to take to get home, and his options are: go straight ahead, turn left, and turn right. Suppose that only one path will lead him home, and that he decides to turn right, because he is most confident that turning right is the correct option for him. Yet intuitively, he might not believe outright of any option that it is the correct one, if his credence in each option being the correct one is less than 0.5. So on the face of it, binary beliefs, unlike credences, cannot explain Jack's decision to turn right. But one might think that binary beliefs are well up to the task. Perhaps Jack's decision is due to his binary belief that turning right has the highest objective probability of being the correct option. On this view, instead of saying that Jack's credence in turning right being the correct option is, for example, 0.49, we might say that Jack has a binary belief that the objective

¹For my purposes, objective probabilities are just non-subjective probabilities. More will be said about the different kinds of objective probability in section 2.3.

probability of turning right being the correct option is 0.49. Call the view that talk of credences is in general replaceable by talk of binary beliefs about objective probabilities the *Replacement Thesis*.

The foregoing view yields a rather unorthodox kind of subjective decision theory, construed as a descriptive enterprise—insofar as the explanation and prediction of action and behaviour is concerned, what matters is one’s set of binary beliefs about the relevant objective probabilities, and not the objective probabilities themselves, or one’s credences. We might also represent the expected utility of an act as $\sum_{i=1} P_o(s_i|a)U(s_i\&a)$, where a is the proposition that one performs the act in question, and $P_o(s_i|a)$ is what is believed to be the objective probability that a certain proposition is true, conditional upon a . For example, to calculate the expected utility of turning right for Jack, we might take the product of what Jack *believes* to be the objective probability of his getting home conditional on his turning right, and his utility in turning right and getting home.² We may refer to such expected utilities as *believed expected utilities*.

One might also continue doing a kind of prescriptive or normative decision theory, by holding that we ought to do whatever maximises our believed expected utilities. Furthermore, one might hold that there are rationality constraints on binary beliefs about objective probabilities that mirror the axioms and theorems of probability theory. For example, if p is a tautology, then instead of holding that S’s credence in p ought to be 1, one might hold that S ought to believe that the objective probability that p is true is 1. Or suppose that p entails q . Instead of holding that S’s credence in q ought not to be less than her credence in p , one might hold that S ought to be such that if she believes that the objective probability that p is true is x , and that the objective

²One might wish to appeal, not to Jack’s utilities, but to what Jack believes to be his utilities. Since my focus is on beliefs and credences, rather than on utilities, I shall set aside this particular way of modifying decision theory.

probability that q is true is y , then $y \geq x$, for all $x \in [0, 1]$ and all $y \in [0, 1]$.

But can binary beliefs really do all the work that credences do? Some philosophers think so. For example, John Pollock (2006) thinks that

once we have decided to give objective probabilities another look, there is an obvious suggestion regarding how to understand these degrees of confidence. *They are just the believed values for the objective probabilities.* You are more confident that p will be true than that $\neg p$ will be true just in case you think p is objectively more probable than $\neg p$. (p. 94; emphasis mine)

Similarly, Gilbert Harman (1986) thinks that ‘whatever principles are developed for changing binary belief will apply to changing degrees of belief, *treating these as binary beliefs about probabilities*’ (p. 24; emphasis mine).

Harman acknowledges that there is such an intuition that beliefs come in varying strengths. But he suggests that

‘these varying strengths [...] are to be explained as a kind of epiphenomenon resulting from the operation of rules of revision. For example, it may be that p is believed more strongly than q if it would be harder to stop believing p than to stop believing q , perhaps because it would require more of a revision of one’s view to stop believing p than to stop believing q . (p. 22)

Harman’s suggestion is problematic. Suppose that a scientist believes fairly strongly that the radioactive atom in the room will decay in the next hour. This might be because she believes *fairly* strongly a scientific theory according to which the chance of the atom decaying in the next hour is quite high. But suppose also that she believes much more strongly that she will not become a millionaire overnight. Given Harman’s suggestion as to how to understand the varying strengths of our beliefs, the scientist will find it much harder to

stop believing that she will not become a millionaire overnight than to stop believing that the radioactive atom will decay in the next hour. According to Harman, giving up the first belief will require more of a revision of her view than giving up the second belief. But this does not seem right. Giving up the belief that she will not become a millionaire overnight might not require a very radical change in her views. It might not be all that hard for her to acquire suddenly the belief, perhaps via a phone call, that a rich relative who is about to die has unexpectedly bequeathed her entire fortune to her. But giving up her belief that the atom in the room will decay in the next hour might require her to give up her pet scientific theory, which she has invested her life defending, as well as other related and established theories that she and other members of her scientific community hold dear. So even though her belief that she will not become a millionaire overnight is stronger than her belief that the radioactive atom will decay in the next hour, it might turn out that giving up the first belief is easier and requires less of a departure from her current views than giving up the second belief.

Also, intuitively, beliefs of the same strength may differ in how easy or hard it is to abandon them. For example, suppose that a coin is about to be tossed, and Jill is as confident that the coin will land heads as she is that it will land tails. If this is because she is certain that the coin is a fair coin, then being told that the coin landed heads in a string of five previous tosses is not going to decrease her credence in the coin landing tails. But if her initial credence in the coin landing tails is due to her having no more reason to think that the coin will land heads than that it will land tails, then upon acquiring the new information, her credence in the coin landing tails might well decrease. The strength of Jill's initial belief that the coin will land tails is the same in both cases, but it is much easier for her to change her opinion about whether

the coin will land tails in the second case than in the first case.³

Finally, as Keith Frankish (2009) has observed, Harman's attempt to explain the intuition that beliefs come in varying strengths does not work when we consider a belief of intermediate strength (p. 78). Intuitively, a person could be somewhat but not very confident that p . To explain this intuition, Harman would have to say that this person has a binary belief that p that is somewhat but not too difficult to give up. But this is implausible. For instance, if the person is only 55% confident that p , it is a stretch to say that she has a binary belief that p that is somewhat but not too difficult to give up—instead, she simply lacks a binary belief that p .

2.2 Problems with Subjective Probabilities?

Harman's attempt at explaining the intuition that beliefs come in varying strengths looks unpromising. But in the first place, why subscribe to the Replacement Thesis? Here's an argument in favour of such a view—a version of the argument can be found in Pollock (2006).

Either subjective probabilities are meant to represent actual credences, or they are meant to represent rational credences. If the first, then we get a rather unrealistic account of the way we actually act. For by definition, subjective probabilities as traditionally understood obey the probability axioms (more of which in the next chapter), and arguably, credences often flout such axioms. For example, our credences often fail to satisfy Normalisation, according to which the subjective probability assigned to any tautology is 1. But if subjective probabilities are meant to represent rational credences, then

³Jill's information that the coin is fair in the first case is not reflected in her credence in the coin landing tails. According to Skyrms (1977), Popper thinks that such extra information can only be reflected in one's credences about objective probabilities. But as Skyrms notes, attributing the point to Jeffrey, such extra information may be reflected in one's conditional credences (p. 707; Jeffrey, 1983, p. 46). Jill's credence in the coin landing tails on the sixth toss, given five previous tosses of heads, remains at 1/2 when she is certain that the coin is fair. But when she's not, such a conditional credence may well be lower than 1/2.

we get a rather unrealistic account of how we ought to act. It seems too demanding to require that we have full credence in any tautology whatsoever on pain of irrationality.⁴ So either way, subjective probabilities, and hence subjective decision theory, do not help provide a good account of how we act or ought to act. The natural replacement for subjective decision theory is a kind of decision theory in which objective probabilities reign. Insofar as we want to explain how we act or ought to act using decision theory, we should appeal to binary beliefs about objective probabilities, rather than to subjective probabilities.

The argument is too quick. Even if subjective decision theory as traditionally understood is flawed, we should be wary about capitulating to a decision theory that eschews credences in favour of binary beliefs about objective probabilities. For given the latter kind of decision theory, we would need to appeal to beliefs about objective probabilities to explain and predict action and behaviour. But then, there will arise similar issues of how to come up with constraints on such beliefs that are not too demanding. For example, it would be too demanding to require that for any tautology, one believes that the objective probability that it is true is 1. And it would also be too demanding to require that for any p and q , if p and q are logically equivalent, then one believes that the objective probability that p is true is x if and only if one believes that the objective probability that q is true is x , for some $x \in [0, 1]$.

Also, subjective decision theory provides us with a formal tool to help explain and predict action and behaviour. If, however, it proves to be inadequate in this respect, the fault lies with the tool, and not with what is modelled. For example, although using the basic normal modal logic **K** to model binary

⁴One might reply that subjective probabilities are meant to model the credences of ideally rational agents. But then, one has to contend with the question of what the credences of ideally rational agents have to do with flesh and blood agents who are not logically omniscient. More on this in the next chapter.

beliefs leads to what many consider to be the unrealistic result that agents believe (in the binary sense) all logical truths, it would be too quick to conclude that there are no such things as binary beliefs. Just as one may attempt to weaken **K** to arrive at a more realistic model of binary beliefs, one may attempt to weaken the constraints that the probability axioms impose upon credences or rational credences. Now Pollock (forthcoming) points out that ‘[g]iven the probability calculus and Bayesian conditionalization, Bayesians purport to prove many interesting and powerful theorems about epistemological concepts like theory confirmation’. The worry is that we weaken the axioms on pain of losing such theorems. In the next chapter, I shall address such a worry; I shall also propose constraints on credences that are more realistic than those provided by the probability axioms.

2.3 All-or-Nothing Beliefs about Objective Probabilities

2.3.1 The Replacement Thesis

For all that has been said so far, one might still think that the Replacement Thesis holds. The rest of this chapter is devoted to showing that the thesis is false.

Suppose I claim that Jack is 60% confident that p . According to Pollock and Harman, I can be understood as claiming that Jack believes that the objective probability that p is true is 0.6. But the latter claim is ambiguous. Do I have in mind a particular concept of objective probability, for instance, objective chance? If so, I might be understood as claiming that Jack believes that the objective chance that p is true is 0.6. But I might have in mind a concept of objective probability without having in mind a concept of any particular kind of objective probability. Perhaps I am just claiming that Jack believes that the *objective probability* that p is true is 0.6, without being committed to

the claim that he believes that the chance (for example) that p is true is 0.6. It is worth drawing this distinction in case an argument that works against understanding a claim like ‘Jack is 60% confident that p ’ in the first way does not work against understanding it in the second way.

The Replacement Thesis admits of various disambiguations depending on how ‘objective probability’ is understood. In what follows, I first shall survey, rather briefly, various kinds of objective probability, including classical probabilities, epistemic probabilities, relative frequencies, and propensities.⁵ I shall then consider and assess various arguments that attempt to show that the Replacement Thesis, on its various disambiguations, is false.

2.3.2 Different Kinds of Probability

The *classical probability* of an event happening is given by the ratio of the number of equipossible outcomes in which the event occurs to the total number of equipossible outcomes there are. Suppose we throw a symmetrically-balanced die. The classical probability that it will land on an even number is $1/3$, calculated by taking the ratio of the number of equipossible outcomes in which the die lands on an even number, namely, 3, to the number of equipossible outcomes in which the die lands, namely, 6. Now, there are issues regarding what ‘equipossible’ means. If saying that two outcomes are equipossible amounts to saying that they are equally probable, then cashing out classical probability in terms of equipossibilities will be circular. The *principle of indifference* says that two possibilities are equipossible when no evidence favours one possibility over another, because we have no evidence whatsoever for any of the possibilities, or because the evidence we have for each possibility is of the same strength.

Epistemic probabilities are conditional probabilities, and are used to

⁵See Hájek (2009) for a detailed discussion of the various interpretations of probability.

capture the intuitive idea that different pieces of evidence confer different degrees of support upon a hypothesis. Intuitively, for example, the epistemic probability that it will rain, given that dark clouds have gathered, is higher than the epistemic probability that it will rain, given that they have not. Different philosophers have different theories of epistemic probabilities and different names for such probabilities. Rudolf Carnap (1950), with his theory of *logical probability* and his notions of state descriptions and confirmation functions, seeks to systematise the idea of a piece of evidence conferring support on a hypothesis by developing a theory of inductive logic. Timothy Williamson (2000) talks about *evidential probabilities*, which are probabilities conditional upon one's total evidence, where one's evidence is just what one knows (pp. 203-207). Mark Norris Lance (1995) talks about *fair rational betting odds* given one's evidence (pp. 22-27). And William Alston (2005), unlike the other philosophers mentioned, thinks that the epistemic probability of a logical truth conditional on some contingent proposition may well fall short of 1 (p. 97-98). One thing worth noting is that just having beliefs about the non-contingent issue of what degree of support one proposition renders another is often not enough to tell us what we ought to do. For epistemic probabilities to help serve as a guide to action, we also need to have an opinion as to what the *relevant* epistemic probabilities are, where an epistemic probability is *relevant* if (like Williamson's evidential probabilities) it is conditional upon our total evidence.

The notions of probability mentioned in the previous two paragraphs have something in common: they involve the concept of evidence. But there are other notions of probability that do not do so. According to frequentism, probabilities are, roughly speaking, *relative frequencies*. According to finite frequentism, the probability of event *A* occurring is the actual relative frequency of events similar to event *A* occurring within a finite reference class. For example, the probability of a smoker dying of lung cancer before age 70 is

the frequency of actual smokers who die of lung cancer before age 70, relative to the number of actual smokers there are. According to hypothetical frequentism, the probability of an event occurring is the limiting relative frequency of the event occurring within a hypothetical, infinite reference class. For example, the probability of a smoker dying of lung cancer before age 70 is the limiting frequency of hypothetical smokers who would die of lung cancer before age 70, relative to the class of infinitely many hypothetical smokers.

According to propensity theories of probability, probabilities are cashed out in terms of *tendencies* or *dispositions* to produce certain outcomes. Suppose that the probability that a particular die will land on an even number is $1/3$. On a long-run propensity theory, this means that the die has a very strong disposition or tendency to land on an even number with a long-run relative frequency of $1/3$. On a single-case propensity theory, the value $1/3$ is a measure of the disposition or the tendency of the die to land on an even number.

2.3.3 Christensen's Argument against the Replacement Thesis

Given the various kinds of objective probability, the Replacement Thesis admits of several disambiguations. For example, one may understand it as the thesis that talk of credences is replaceable by talk of binary beliefs about chances, or about epistemic probabilities. But David Christensen (2004) has argued that the Replacement Thesis risks 'attributing to the agent a belief about matters too far removed from the apparent subject matter of her belief' (p. 19). Suppose that a fan of the Replacement Thesis claims that 'Ida is 60% confident that she will not live to a hundred' can be understood as something like 'Ida has a binary belief that the relative frequency of people like her living to a hundred is 0.6'. The worry is that the concept of the particular kind of objective probability involved might simply be too sophisticated

for an agent to possess.⁶ For example, Christensen notes that if the relevant objective probability is relative frequency, then the agent needs to possess (at least implicitly) a concept of reference classes. And if the relevant objective probability is propensity, then the agent needs to possess (at least implicitly) a concept of disposition. But arguably, such concepts are too sophisticated for some agents to possess, even if we are ordinary happy to attribute credences to them. Likewise, Keith Frankish (2009) thinks that credences are not binary beliefs about epistemic probabilities, or about behavioural dispositions to bet, since children and animals do not seem to possess the concepts of evidential relations and betting odds (*ibid.*).

Patrick Maher (2006) is no fan of the Replacement Thesis, but he is not convinced by arguments like Christensen's or Frankish's (pp. 134-135). He thinks that the concepts of *physical probability* and *inductive probability* are ordinary and vague, and do not require a great deal of conceptual sophistication to possess. To illustrate these concepts, he considers a case in which a coin that is either two-headed or two-tailed—we do not know which—is tossed. The physical probability that it will land heads is either 0 or 1, but the inductive probability that it will land heads, given one's relevant evidence about the coin, is 1/2. According to Maher, we can make such probability judgements even if we do not possess the concepts of reference classes and evidential relations. He thinks that the more sophisticated concepts of frequency and propensity are really explicata of the everyday and less sophisticated concept of physical probability. And although he does not say so explicitly, it seems that he would take Carnap's concept of logical probability to be an explicatum of inductive probability. If Maher is right, a proponent of the Replacement Thesis may

⁶Another related worry suggested by the quote from Christensen seems to be that even if the agent is sophisticated enough to possess the relevant concepts, the Replacement Thesis changes the subject matter of Ida's belief. Her belief is about her living to a hundred years old, and *not* about the relative frequency of people like her living to a hundred years old.

still hold that we can understand a claim like ‘Jack is 60% confident that p ’ as ‘Jack believes that the physical probability (or the inductive probability) that p is true is 0.6’.

Suppose we grant that even the most naive among adult human beings have the concepts of physical probability and inductive probability. That still leaves open the issue of whether very young children or animals have such concepts. Perhaps to have such concepts, one needs to have the concept *physical* and the concept *evidence*, which arguably, very young children or animals lack. One may, however, go a step farther than Maher, and claim that it is possible to have an ordinary and vague concept of probability, without having a concept of either physical or inductive probability in mind. Perhaps when one thinks, ‘It is likely that it’ll rain this afternoon’, one is attributing a high probability to some event happening, even if one does not know anything about what evidence is or about what ‘physical’ means. On this view, a proponent of the Replacement Thesis might understand ‘Jack is 60% confident that p ’ as ‘Jack believes that the probability that p is true is 0.6’, where ‘probability’ is some ordinary and vague concept that even children and animals might have.

However, do children and animals possess even this ordinary and vague concept of probability? Richard Jeffrey (1983) thinks that cats have credences; how plausible is it to think that such credences are really binary beliefs about probabilities (p. 70)? Jeffrey himself holds that when we attribute a propositional attitude to a cat, we need not imply that the cat is thinking at all. He holds that the ‘theory of deliberate action is ours, not the cat’s, and the theory can be used to explain some of the cat’s actions even though the cat does not understand the theory, just as the cat can digest his food without being a chemist’ (*ibid.*). Perhaps if attributing to the cat beliefs about probabilities helps make sense of the cat’s behaviour, then it is not too silly to do so.

We cannot really settle the issue of whether cats and children have a

concept of probability until we have a firmer grasp of what it is to possess a concept. One may think that cats do not really have any beliefs, and that we are anthropomorphising when we attribute credences to them.⁷ But even though Christensen's argument against the Replacement Thesis leaves some wiggle room for its proponent, it at least shows what the latter is committed to. Insofar as she is willing to attribute doxastic states to cats, she will have to say that cats possess some concept of probability. To the extent that she finds this implausible, she should rethink her endorsement of the Replacement Thesis.

2.3.4 Other Problems for the Replacement Thesis

The Replacement Thesis faces other problems. There are cases in which one's credence in p is x , for some $x \in [0, 1]$, but one does not believe that the objective probability that p is true is x . This might happen in at least two kinds of case. First, one believes that it does not make sense to speak of *the* objective probability that p is true. And second, one believes that the objective probability that p is true is not x . As we shall see, one way the first might happen is when one has an unconditional credence in p , but one is confident that the objective probability in question is either conditional in nature or relative to some restricted domain. And one way the second might happen is when one's credence in p corresponds to a mixing of objective probabilities.

Consider a one-off event that has already happened, say the 2008 US presidential election. And suppose that Ida, who was in a coma during the entire election, has just awoken, and has a credence of 0.8 that Obama won. According to the Replacement Thesis, Ida believes that the objective probability that Obama won is 0.8. Suppose that the relevant objective probability is actual relative frequency. Then since the event concerned is a one-off event,

⁷For more on the attribution of beliefs to animals, see Stalnaker (1984).

and there are no events like it, the relative frequency of its occurrence is either 0 or 1. But if Ida is confident of this, then contrary to the Replacement Thesis, she does not believe that the objective probability that Obama won is 0.8. Now suppose that the kind of objective probability at issue is hypothetical limiting relative frequency. To compute the relative frequency, we would have to consider a reference class containing hypothetical events like the 2008 US presidential election, in which people like Obama stand for election. But there is a worry having to do with reference classes. When one inquires as to the relative frequency of an event, one has to deal with the question, ‘Relative to which reference class?’ For there is no such thing as *the* reference class to which the event belongs.⁸ Which reference class we appeal to in computing relative frequencies may ultimately depend on pragmatic considerations, and Ida may believe that the relative frequency of Obama winning the election is 0.8, relative to a particular reference class, and that it is 0.85, relative to another. But if her credence that Obama won the election is 0.8, period, then a proponent of the Replacement Thesis will have to say that there is a reference class that is privileged over all others.

Appealing to propensities instead of relative frequencies is not going to help. Since the election has already happened, it might not make sense to speak of how there is currently a disposition or tendency for Obama to win the election in 2008—he has already won it. So Ida might think that it does not make sense to speak of the single case propensity of Obama winning the election. Or even if she thinks that it makes sense to speak of it, she might think that the propensity is either 0 or 1. Appealing to long-run propensities is not going to help either. In fact, an appeal to long-run propensities will bring in its train the problems faced by appealing to relative frequencies, since long-run propensities are closely related to relative frequencies. Ida may believe

⁸See Hájek (1997).

that the long-run propensity of Obama winning the election is 0.8, relative to a particular reference class, and that it is 0.85, relative to another. But if her credence that Obama won the election is 0.8, period, then a proponent of the Replacement Thesis will have to say that there is a reference class that is privileged over all others.

As for classical probability, there will be cases in which one is confident that there is no such thing as *the* classical probability that p is true. It is well-known that when the outcomes in question can be partitioned in uncountably many ways, applying the principle of indifference to determine equipossible outcomes leads to paradox.⁹ One might hold that it makes sense to speak of classical probability only in cases in which the outcomes in question cannot be partitioned in uncountably many ways. But even something as ordinary as the throw of a die might have uncountably many possible outcomes—the die might land on an even number in a particular way, at a certain angle, at a certain distance from the edge of the table, etc. So it looks like the Replacement Thesis is severely limited if the kind of probability involved is classical probability.

Similarly, in many cases, even though one has a sharp credence in p , one might be confident that there is no such thing as *the* epistemic probability that p is true. Or at least, the relevant epistemic probability might range over an interval of values, instead of taking a sharp value. Again, such cases might be ones in which we are led to paradox if we apply the principle of indifference, or more generally, ones in which our evidence does not determine a uniquely

⁹The following example is taken from Hájek (2009), who adapted it from van Fraassen (1989). If all we know about a cube factory is that it produces cubes with side-lengths between 0 and 1 metre, what is the probability that a random cube from the factory has a side-length between 0 and 1/2 a metre? It is natural to think that the evidence for the cube having a side-length in the interval $[0, 1/2]$ is the same as the evidence for the cube having a side-length in the interval $[1/2, 1]$. Applying the principle of indifference in such a case yields an answer of 1/2. Now equivalently, what is the probability that the face-area is between 0 and 1/4 of a metre square? It is natural to think that the evidence for the cube having a face-area in the interval $[0, 1/4]$ is the same as that for the cube having a face-area in the intervals $[1/4, 1/2]$, $[1/2, 3/4]$ and $[3/4, 1]$. Applying the principle of indifference in such a case yields an answer of 1/4. Contradiction.

relevant epistemic probability.¹⁰ For example, one might think that if one has no evidence whatsoever about the dimensions of cubes that are made in a cube factory, then the relevant epistemic probability that a random cube made in the factory has a side-length between 0 and 1/2 a metre should not be 1/2; rather it should have the imprecise value $[0,1]$. One might hold that it is irrational to have a sharp credence in p when our evidence does not determine a uniquely relevant epistemic probability that p . But the Replacement Thesis purports to replace credences, and not just rational credences, with binary beliefs about objective probabilities.

Let us discuss in greater detail the kind of case in which one's credence in p is x , but one believes that the objective probability that p is true is not x . Consider the earlier example from Maher that involves tossing a coin that is either two-headed or two-tailed, but we do not know which. We might well believe that the propensity or the relative frequency of the coin landing heads is either 0 or 1. But it seems reasonable for our credence in the coin landing heads to be 1/2, corresponding to mixing the two possible values, 0 and 1, for the objective probability that the coin will land heads.¹¹ Now it might be thought that the relevant kind of objective probability in this case is epistemic probability or fair rational betting odds. But consider another example from Maher that involves a statement h that either entails e or $\neg e$, but we do not know which. As Maher points out, we might well believe that the epistemic probability (or Maher's inductive probability) that h entails e is either 0 or 1. But our credence in h entailing e might well be 1/2, corresponding to mixing the two possible values, 0 and 1, for the relevant epistemic probability that h entails e . Similarly for fair rational betting odds on h entailing e .

One might think that all Maher's example shows is that we need a no-

¹⁰See Hájek and Smithson (forthcoming) and Peter Walley (1981) for more on imprecise and/or indeterminate probabilities.

¹¹As Hájek (ms.) points out.

tion of epistemic probability that is suitable for non-ideal creatures like human beings, and that takes into account our mathematical or logical ignorance. Perhaps there is nothing in the ordinary and vague concept of inductive probability that requires that mathematical and logical truths get assigned a probability of 1. If anything, it is natural to think that ordinary people in everyday situations sometimes have non-conclusive evidence that certain mathematical or logical statements are true. For example, take testimonial evidence for the truth of a particular logical statement. If a logic enthusiast tells you that h entails e , you might take that as evidence that h entails e . There might be some notion of epistemic probability such that the epistemic probability that h entails e , conditional on such evidence, has an intermediate value between 0 and 1. If this is right, then to explicate Maher's notion of inductive probability, it is not enough to appeal to epistemic probability as traditionally conceived. We will need to invoke some notion of *non-ideal* epistemic probability.

In fact, Alston (2005) has argued for a theory of epistemic probability according to which we need not assign epistemic probabilities of 1 to mathematical or logical truths (pp. 97-98.) Suppose that such a theory works. Consider again Maher's example in which a statement h either entails e or entails $\neg e$, but we do not know which. On a theory of non-ideal epistemic probability, it might turn out that the epistemic probability of h entailing e , given one's evidence, is $1/2$. In such a case, instead of saying that our credence in h entailing e is $1/2$, one could perhaps say that we believe that the non-ideal epistemic probability of h entailing e is $1/2$. Similarly, we might be able to understand fair rational betting odds in a way such that the fair rational betting odds on h entailing e are 1:1.

Yet it is scant consolation for the fan of the Replacement Thesis even if there is some theory of non-ideal epistemic probability that works. For there might be cases in which one is not sure what one's total evidence is. Suppose

that one is unsure whether one's total evidence is e_1 or e_2 , although one is certain that it is one of the two. Suppose also that as a result, one is not sure whether the relevant (possibly non-ideal) epistemic probability that h is true is 0.85 or 0.9 (for example), although one is sure that it is one of the two. One could well have a credence in h that has a value that is neither 0.85 nor 0.9. But by hypothesis, one does not believe that the relevant epistemic probability that h is true has such a value. One way to deny that such a case is possible is to deny that one can be uncertain of what the relevant epistemic probability is. But this is implausible—a non-ideal agent need not always have access to the relevant epistemic probabilities, be they ideal or non-ideal. Similarly, we may not be sure what the fair rational betting odds on the truth of p are. Hence, what we believe to be the fair rational betting odds on the truth of p need not match our credence in p .

Might there be a new kind of objective probability that has not occurred to us, and that avoids the problems that the Replacement Thesis faces? Or can we avoid the problems that the foregoing arguments pose for the Replacement Thesis by holding that in formulating the thesis, all we need to do is to appeal to some ordinary and vague concept of objective probability, without appealing explicitly to any particular kind of objective probability?

Here's an argument that shows why the answer to both questions is 'No'. There are philosophers, e.g. Jeffrey (1997), who think that there are no such things as objective probabilities (of any kind), or at least, that if there are objective probabilities, they are to be understood as some sort of robust credences. Suppose Jeffrey's credence in p is x . But if he does not believe that there are such things as objective probabilities (other than robust credences), then *a fortiori*, he does not believe that the objective probability that p is true is x . Fans of the Replacement Thesis will have to deny that such a case is possible. They will have to maintain that Jeffrey believes that the objective

probability that p is true is x , and hence, believes that objective probabilities exist. They will have to hold that if Jeffrey sincerely insists that he does not believe that objective probabilities exist, then it just shows that he has a false belief about his own beliefs. Now I can't prove with ironclad certainty that such a reply is incorrect, but I hope its implausibility speaks for itself.

Although problems abound for the Replacement Thesis, one might have to live with them if there are even more severe problems with credences. As we saw in section 2.2, Pollock worries that if credences or rational credences obey the probability axioms, then we get an unrealistic account of how we act or ought to act, but if they don't, then we lose certain important results in Bayesian epistemology. The next chapter addresses this worry.

Chapter 3

Credences and Logic

3.1 The Probability Axioms

Call the view that rational credences obey the standard probability axioms *probabilism*. A charge frequently levelled against probabilism is that it makes the unrealistic demand that rational agents be logically omniscient, assigning a credence of 1 to any logical truth, and a credence of 0 to any logical falsehood. As mentioned in chapter 1, the supposed fecundity of subjective decision theory gives us reason to hold that credences rather than binary beliefs do all the heavy lifting in the explanation and prediction of action and behaviour. But such reason looks less compelling if subjective decision theory requires probabilism to be true, and if probabilism demands logical omniscience. In fact, as we saw in chapter 2, Pollock (2006; forthcoming) holds that the falsity of probabilism is reason to embrace binary beliefs instead of credences.

In this chapter, I shall explain why probabilism is thought to lead to logical omniscience—strictly speaking, it doesn't—and discuss some of the problems that probabilism faces. After that, I shall propose and defend certain constraints on rational credences that improve upon those provided by the probabilist. These constraints will allow us to deal with agents like human beings who fall short of logical omniscience, but whom we would still want to

deem more or less rational. Since the constraints are supposed to be constraints on rational credences, they are constraints that the credences of such agents ought to satisfy. But since such agents are more or less rational, the constraints will also help provide a good approximate description of how such agents will behave. So although I focus on rational credences, the constraints I defend are intended to be more descriptively accurate than those provided by the probability axioms—even if our credences do not conform to the constraints perfectly well, they conform to them better than they conform to the axioms.

Let's begin by looking at the axioms of standard probability theory. Let probabilities be assigned to declarative sentences in some formal language \mathcal{L} , say, the language of standard propositional logic. Let the sentences in \mathcal{L} be closed under the connectives $\neg, \vee, \wedge, \equiv$ and \supset —we can derive the latter three connectives from the former two in the usual way. And let P be a probability function from \mathcal{L} to the real-valued interval $[0,1]$ satisfying the following axioms:

- a. (Non-negativity) $P(p) \geq 0$, for all $p \in \mathcal{L}$.
- b. (Normalisation) If t is a tautology, then $P(t) = 1$, for all $t \in \mathcal{L}$.
- c. (Finite Additivity) $P(p \vee q) = P(p) + P(q)$, for all $p \in \mathcal{L}$ and $q \in \mathcal{L}$ such that p and q are logically incompatible.¹

The following is often taken to be the definition of conditional probability:

- e. (CP) $P(p|q) = \frac{P(p \wedge q)}{P(q)}$ if $P(q) \neq 0$, for all $p \in \mathcal{L}$ and $q \in \mathcal{L}$.

And from the axioms, we can derive the following two theorems:

- g. (Entailment) If p entails q , then $P(q) \geq P(p)$, for all $p \in \mathcal{L}$ and $q \in \mathcal{L}$.

¹I shall set aside the issue of whether Countable Additivity should be a rationality constraint on credences, since even probabilists disagree about the issue. For example, see Bartha (2004). In any case, since Countable Additivity entails Finite Additivity, the former should not be a rationality constraint on credences if the latter should not.

- f. (Equivalence) If p and q are logically equivalent, then $P(p) = P(q)$, for all $p \in \mathcal{L}$ and $q \in \mathcal{L}$.

In themselves, the probability axioms, elegant as they are, do not tell us anything about rational credences. We have a set of mathematical apparatus, but it is sterile unless we breathe philosophical life into it. We do this by deciding what P and the sentences in \mathcal{L} represent. Suppose we let P_j represent Jack's rational credence function at a particular time, and we let the sentences in \mathcal{L}_j represent the propositions to which Jack's credences are assigned. Then $P_j(p)$ represents Jack's rational credence in p at that time.

We can see why there might be a problem of logical omniscience. By Normalisation, $P_j(t) = 1$, for any $t \in \mathcal{L}_j$ that is a tautology. But it seems too demanding to hold that human beings who assign credences of less than 1 to complex logical truths are irrational. Perhaps the kind of rationality at issue is ideal rationality. But then, there is the further question of why we should be interested in or should only be interested in ideal rationality.

I shall deal with two preliminary matters before we investigate the above issues in greater detail. The first has to do with whether to take probabilities to be assigned to sentences or to sets. So far, I have only considered a sentence-based formulation of the probability axioms, even though we could have formulated the axioms in a set-theoretic fashion.² Now if we are dealing with logically omniscient agents, it does not matter much whether probabilities are assigned to sets or to sentences. But some, like Ian Hacking (1967), worry

²Here are the probability axioms formulated set-theoretically. Let Ω be a non-empty set and \mathcal{F} be a set of subsets of Ω that contains Ω as a member, and is closed under complementation and union. Then a probability function, P , is a function from \mathcal{F} to the real numbers satisfying the following axioms:

- a' $P(A) \geq 0$, for all $A \in \mathcal{F}$.
- b' $P(\Omega) = 1$.
- c' $P(A \cup B) = P(A) + P(B)$ for all $A, B \in \mathcal{F}$ such that $A \cap B = \emptyset$.

that if we want constraints on credences that are more realistic than those provided by the standard probability axioms, then adopting a set-theoretic approach will not serve us well (p. 318). For the elements of the sets to which probabilities are assigned are typically taken to be possible worlds, and logical truths are true at all worlds.

Perhaps we can get around the problem by appealing to ‘impossible possible worlds’ at which logical truths are false (Jaakko Hintikka, 1975, p. 477). I am open to such a move, although more has to be said about these impossible worlds to assuage the worries of philosophers who find such worlds suspect.³ To avoid having to deal with such worries, I shall take probabilities to be assigned to sentences rather than to sets. Doing so will also be convenient, since most of the philosophers I discuss in this chapter take probabilities to be assigned to sentences. Note, however, that just because I take probabilities to be assigned to sentences does not mean that I take the objects of credences to be sentences. It is up to us how we interpret probability assignments—as mentioned earlier, it is up to us what we take P and the sentences in \mathcal{L} to represent. For example, let the atomic sentences in \mathcal{L}_j be English sentences. I may interpret ‘ $P_j(\text{It is raining}) = 0.8$ ’ as saying that Jack’s credence that it is raining is 0.8, even if Jack knows no English.

The second preliminary matter concerns certain views about belief content.⁴ Some philosophers, including Ruth Barcan Marcus (1983), Robert Stalnaker (1984), and Lewis (1986), hold that belief content is to be cashed out in terms of possible worlds (or centred possible worlds, in the case of Lewis), where by *possible worlds*, they mean to exclude impossible worlds. So they are happy to embrace the view that we all satisfy Normalisation—even cats and

³For example, see Lewis (1998), who wonders how impossible worlds are constructed, and Stalnaker (2003), who thinks that impossible worlds are ‘too much to swallow’ (p. 123; p. 55). See Nolan (1997) for an attempt to deal with worries about impossible worlds.

⁴I am indebted to Jens Christian Bjerring for discussion on several of the issues raised in the remainder of this section.

dogs, if they have beliefs. But on such a view, there is just one necessarily true proposition (albeit expressible in a myriad of ways), and any agent, whether ideally rational or not, cannot help but have a credence of 1 in it. In such a case, someone like Pollock can't point to how Normalisation makes unrealistic demands on *rationality* as a reason to eschew credences in favour of binary beliefs.

Of course, even Stalnaker will not deny that we often act in ways that *seem* to indicate that we do not have a credence of 1 in the necessary proposition. But according to Stalnaker (1984), when we *apparently* fail to believe what is necessarily true, what we really fail to believe is the *contingent* and *metalinguistic* proposition that some sentence, *s*, standardly used to express the necessary proposition, expresses the necessary proposition (pp. 84-87). In such a case, we might well not be disposed to assent to the sentence *s*. Or we might even be disposed to bet that the sentence *s* expresses a falsehood if the payoff in the event of a win is sufficiently attractive. And betting against *s* or failing to assent to *s* may be a sign of epistemic irrationality if we understand in different fragments of our mind the components that make up *s*, but we somehow can't put the fragments together to see that *s* is standardly used to express the necessary proposition.⁵ Now can't we take Normalisation to say that for any sentence that expresses a necessary proposition and whose components an agent understands, the agent ought to have a credence of 1 in the contingent proposition that the sentence expresses the necessary proposition?

Perhaps we can re-interpret Normalisation in such a fashion, but it will run afoul of any requirement that credence functions be *regular*, where a credence function is regular if and only if no contingent proposition is assigned a credence of 0 or 1. If you think that credence functions ought to be regular, or

⁵See Stalnaker (1984), pp. 84-87. If we don't understand the components that make up *s*, then our betting against the sentence or disposition not to assent to it might be a sign of mere linguistic ignorance, rather than a sign of irrationality *per se*.

at least, that we ought not assign credences of 0 or 1 to contingent propositions about whether certain sentences express the necessary proposition, you will balk at this way of thinking about probabilism. In any case, for the rest of this chapter, I shall speak as if one can fail to have full credence in necessary truths. If you harbour views like Stalnaker's, and you're happy with the assignment of credences of 0 or 1 to certain metalinguistic propositions, reinterpret what I say accordingly.

3.2 Problems with Probabilism

The probability axioms are typically understood to be necessary but insufficient constraints on rational credences.⁶ One may well obey the axioms, but fail to satisfy other requirements of rationality. For example, one may fail to satisfy Lewis's Principal Principle (Lewis, 1980). Note also that the kind of rationality at stake is epistemic, rather than pragmatic in nature. According to the probabilist, it is epistemically irrational to have a credence of less than 1 in 'Either it is raining or it is not'. But it might be pragmatically rational to try to put oneself in such a state, if the alternative is eternal torture.

Further, let us distinguish between a strong and a weak version of probabilism. On the strong version, one's credences have to obey the axioms on pain of irrationality of an epistemically blameworthy kind. On the weak version, which is by far the version that is more commonly subscribed to, one's credences have to obey the axioms on pain of falling short of *ideal* rationality. This distinction parallels the familiar distinction in ethics between failing to perform morally obligatory acts, and failing to perform supererogatory acts. The first makes one morally reprehensible; the second just means that one has fallen short of sainthood.

⁶But some philosophers, notably de Finetti (1937), think one is rational so long as one's credences obey the axioms, and one updates one's credences by conditioning.

3.2.1 Too Demanding

The problem with strong probabilism is that it is too demanding. Due to our limited cognitive abilities, there will be logical truths so complex that even the best logicians will not be able to tell that they are true. Furthermore, there are cases in which we would be irrational to assign credences of 1 to logical truths. Suppose there is a complex logical truth to which you assign a high credence only because of the testimony of a logician you are confident is highly reliable, but ultimately fallible. Since the logician is highly reliable, it seems rational for you to have a high credence in the logical truth. But it seems irrational for you to have full credence in it, since you are confident that the logician is fallible.

The charge of being too demanding is less damaging for weak probabilism, since it makes the lesser claim that someone who violates the probability axioms is not ideally rational. But weak probabilism requires that an ideally rational agent be logically omniscient. Now one might think that an ideal should be achievable, or even if unachievable, not too remotely out of reach. But logical omniscience is more than just unachievable; we cannot even hope to get close to it. Supposing that an ideal is something that we can at least get close to, one might wonder if the assumption of logical omniscience is really a mere *idealisation*, rather than an *ideal*. As Hájek (2006) points out, the words ‘idealisation’ and ‘ideal’ are almost homophonous, but although idealisations can be very useful, they carry no normative force by themselves (p. 43).

It is, however, debatable whether an ideal must always be something that we can strive towards or approximate. Christensen (2004) argues that although it is ‘deeply impossible’ to make cars that are perfectly efficient, perfect efficiency nonetheless remains an ideal: even if we have the most efficient car that we can make given the laws of nature, it still makes sense to think that if

the laws of nature had been different, the car could have been more efficient, and hence better in some sense (pp. 166-167). Similarly, a proponent of weak probabilism may hold that even though logical omniscience is unachievable, it nevertheless remains an ideal: even if we have an agent who is as rational as humanly possible, she could have been more rational, and hence better in some sense.

Let's grant that logical omniscience is an ideal, or at least, a useful idealisation. In fact, in discussing the relationship between binary belief and credence, I will sometimes assume that agents are logically omniscient. Still, depending on the job at hand, we may search for constraints on credences that are easier to satisfy. For example, if we wish to explain and predict the action and behaviour of beings like ourselves, who are more or less rational, but not ideally rational, it is reasonable to demand constraints on rational credences that, in the words of Pollock (forthcoming), apply to 'real cognizers with finite reasoning powers, limited memory capacity, etc'. Weak probabilism by itself does not satisfy such a demand.

The above objections against probabilism gesture at the point that probabilism leaves little room for logical inquiry as a rational enterprise. Now I do not wish to deny that, all things being equal, there is something worse about failing to have full credence in a logical truth than failing to have full credence in a contingent truth—in fact, those who think that credence functions ought to be regular might think that it is irrational to have a credence of 1 in a contingent truth. But probability, after all, is the study of uncertainty, and we can be uncertain of logical truths, just as we can be uncertain of many contingent truths. In fact, I may be rational in being less confident that certain complex logical truths are true than I am that certain contingent propositions such as 'I've hands' are true. Given all this, we should take the enterprise of logical inquiry seriously, and let such seriousness be reflected in the constraints

imposed on rational credences.

One might think that probabilism appears to be lacking only because of the failure to distinguish between what Daniel Garber (1983) calls *global Bayesianism* and *local Bayesianism* or what I shall call *global probabilism* and *local probabilism*. Garber made the distinction in dealing with the problem of old evidence, which he thinks is a problem only because of the assumption that agents are logically omniscient. But he thinks that we can avoid making such an assumption if instead of subscribing to global probabilism, we subscribe to local probabilism. The former, according to Garber, deals with ‘the *ideal language of science*, a maximally fine-grained language [...] capable of expressing all possible hypotheses, all possible evidence’ (p. 110; Garber’s emphasis). The latter, however, ‘deal[s] only with the sentences and degrees of belief that [scientists or decision makers] are actually concerned with, those that pertain to the problem at hand’ (p. 111). Garber points out that both global and local probabilism are consistent with the requirement that credences obey the probability axioms—if global probabilism is problematic, the solution is not to give up probabilism wholesale but to adopt the local version of it. Garber’s idea involves having probability functions that are defined over a limited set of sentences, \mathbf{L}^* , with the size of the set and the sentences in it depending on the particular situation that we are trying to study. Some logical truths in such a set of sentences will be treated as unstructured atomic sentences, in which an agent might assign a probability of less than 1. Local probabilism, whether strong or weak, allows us to avoid talking about agents endowed with full blown logical omniscience.

Garber is right that in requiring that credences satisfy the probability axioms, global probabilism is not forced upon us. But as he acknowledges, even though local probabilism avoids full blown logical omniscience, agents are ‘*still* required to know all tautologies of \mathbf{L}^* ’ (p. 117; Garber’s emphasis).

So the constraints on credences imposed by Garber's account are still rather unrealistic. And as Ellery Eells (1990) points out:

There will always be extremely complex logically true sentences of the local language, and extremely simple logically true sentences 'outside' the local language, where it will be inappropriate to insist on probability 1 for the former while not so insisting in the case of the latter. (p. 214)

Here's another worry. The fewer the atomic sentences we allow in \mathbf{L}^* , the more impoverished \mathbf{L}^* will be when it comes to representing the doxastic state of an agent. I grant that there might be propositions that we are incapable of grasping, and to which we do not assign any credence. So there might be good reason not to take the set of sentences over which probabilities are defined to be too fine-grained. But if we want a model of credences that captures an agent's full range of opinions adequately, and that tells us how a person is disposed to act in various actual or hypothetical circumstances, then there is danger in having a very small \mathbf{L}^* . Jack may have an opinion about whether Canberra is the capital of Australia, even if the truth of his opinion is currently irrelevant to the decision problem he is trying to solve, and so 'Canberra is the capital of Australia', according to Garber, is not in \mathbf{L}^* . But then, if P_J represents Jack's credence function, $P_J(\text{Canberra is the capital of Australia})$ will be undefined, and we will be unable to capture Jack's doxastic state adequately. However, the more sentences we allow in \mathbf{L}^* , the more tautologies we can form out of these sentences, and in which the agent ought to assign a probability of 1. On Garber's account, there seems to be an unfortunate trade-off between capturing an agent's full range of opinions adequately, and making realistic demands on her rational credences.

One might think that to avoid this trade-off, we can appeal to different

probability functions defined over different sets of sentences to model Jack’s full range of opinions at a particular time—even if \mathbf{L}^* does not contain ‘Canberra is the capital of Australia’ as an atomic sentence, there is some other set \mathbf{L}^+ that contains it, and that might not contain some of the atomic sentences that \mathbf{L}^* contains.⁷ Each set might just be small enough so that the demands that the probability axioms impose upon the corresponding credences relative to each set are not too unrealistic.

There are two problems with such a move. First, it assumes that the question ‘How confident is Jack that Canberra is the capital of Australia?’ is ill-formed—it assumes, for example, that Jack might at the same instant be very confident that Canberra is the capital of Australia relative to \mathbf{L}^+ , and have no opinion about the matter whatsoever relative to \mathbf{L}^* . But the question does not seem ill-formed. Even when Jack is sleeping or thinking about something unrelated to Canberra or Australia, it might well be correct to say that he is very confident that Canberra is the capital of Australia, full stop. Now one might hold that it is correct to say such a thing so long as $P_J(\text{Canberra is the capital of Australia})$ is defined relative to *some* set of sentences. But this leads us to a second problem with the reply.

Grant that Jack has an opinion towards p so long as $P_J(p)$ is defined relative to some set of sentences. Consider a set $\mathbf{L}^\#$ that contains all the atomic sentences in each of \mathbf{L}^* and in \mathbf{L}^+ . If Jack has opinions about the truth of each atomic sentence in \mathbf{L}^* and he has opinions about the truth of each atomic sentence in \mathbf{L}^+ , then he will have opinions about the truth of each atomic sentence in $\mathbf{L}^\#$. So now there is a question of whether Jack’s credences relative to $\mathbf{L}^\#$ should obey the probability axioms. If $\mathbf{L}^\#$ is big enough, requiring that Jack’s credences obey the axioms might risk imposing demands on them that are too unrealistic. But if Jack’s credences relative to

⁷Thanks to John Cusbert for raising this point.

$L^\#$ are not required to satisfy the axioms, then there is a question of what constraints they ought to satisfy.

3.2.2 Too Beholden to Logic

Another objection less often made against probabilism (in any of its versions) is that it assumes that there is a correct logic, and that we are rationally required to subscribe to it. But consider the fact that philosophers are unsure which logic is the correct logic. Although all claims of the form $p \vee \neg p$ are tautologous in classical logic, this is not so in intuitionistic logic. And although all claims of the form $p \wedge \neg p$ are logically false in classical logic, this is not so in paraconsistent logic. Suppose that probabilists take the truths of logic to be given by classical logic. They are then committed to the view that anyone who doubts the correctness of classical logic, and assigns a credence of less than 1 to propositions that are logically true by the lights of classical logic, is irrational.⁸ This is implausible. Much as one might disagree with philosophers who doubt the correctness of classical logic, one should not deem them irrational in having such doubts. (The objection applies, *mutatis mutandis*, to versions of probabilism that take the truths of logic to be given by some non-classical logic.)

In fact, some may even doubt that there is such a thing as *the* correct logic. For example, following J.C. Beall and Greg Restall (2005), one may be a logical pluralist. Hartry Field (2009), who goes through various versions of logical pluralism, concludes that there is no version that is both true and broad enough to be of significant interest. But so long as one thinks that neither Beall and Restall nor Field are downright irrational, one may be rational in assigning some credence to logical pluralism being true, in which case one might not be irrational in assigning positive credence to some claim of the form $p \wedge \neg p$, even

⁸See Talbott (2008) for a similar objection against probabilism.

if classical logic is in fact the correct logic.

Now the weak probabilist may point out that ideally rational agents would subscribe to the correct logic. Again, there may be no such thing as the correct logic. But even assuming that there is one, weak probabilism offers little help when it comes to searching for constraints on credences that non-ideal agents can hope to satisfy.

3.2.3 Does Scant Justice to Coherence

Probabilism in all its versions also mishandles a certain aspect of rationality that one expects it to handle well, namely, that related to the *coherence* of our credences, of how well they fit together. The term *probabilistically coherent* is a term of art used to describe credences that obey the probability axioms. But sometimes, the flaws of those who fail to satisfy the axioms may well be *logical* rather than *probabilistic* flaws: witness someone who fails to satisfy Finite Additivity only because she fails to be certain that p and q are logically inconsistent.

Here's an intuitive sense in which credences fail to fit together. Consider Jill, who has full credence that p if and only if q , but has different credences in p and in q . Even if it is in fact false that p if and only if q , Jill seems irrational—her credences do not seem to fit well together. Given that Jill satisfies both Non-negativity and Normalisation, the probabilist will hold that the root of Jill's irrationality lies in the violation of Finite Additivity.⁹ But it seems better to just say that there is something *fundamentally* irrational in having full credence that p if and only if q , and yet assigning p and q different credences. The probabilist's explanation of Jill's irrationality locates its source at the wrong place, since it hardly seems to be due to a violation of

⁹It is a probability theorem that if $P(p \equiv q) = 1$, then $P(p) = P(q)$. To get an intuitive grasp of why Finite Additivity has to be invoked in the derivation of this theorem, note that the theorem involves two propositions, and Finite Additivity is the only axiom that does so.

Finite Additivity *per se*.

To illustrate the previous point, suppose we decide that Finite Additivity should not be a constraint on our credences.¹⁰ Even though we would still want to maintain that Jill is irrational, we would no longer be able to appeal to the probabilist's explanation. Note that this problem may actually be more acute for the weak probabilist than for the strong probabilist. Jill, it seems, is patently irrational in an epistemically blameworthy way. Yet all that weak probabilism entitles one to say is that Jill has fallen short of ideal rationality. Admittedly, this need not be a problem for weak probabilism *per se*, as the weak probabilist only aims to offer constraints on ideally rational credences. Why should we not be satisfied with just weak probabilism, and forget the search for constraints on credences that agents like us can reasonably be expected to obey?

3.2.4 Does Not Provide a Good Account of Action and Behaviour

An important reason for not being satisfied with weak probabilism alone lies in the fact that armed only with the constraints it imposes on ideally rational credences, we are unable to explain the action and behaviour of agents who are not logically omniscient, but who are otherwise more or less rational. Consider one such agent who assigns a probability of less than 1 to t , where t is a tautology in the language over which her probabilities are defined. Consider another such agent who assigns higher probability to p than to q , even though p and q are logically equivalent sentences in the language over which her probabilities are defined. Intuitively, the first agent will not bet the farm on t , and the second agent will bet more money on p than on q . (Assume that the agents care only about money, and that the act of betting is independent

¹⁰As indeed, some have argued; for a brief discussion of theories of probability that give up Finite Additivity, see Hájek (2001).

of the truth of p and of q .) Weak probabilism (or strong probabilism, for that matter) is unable to predict or explain the behaviour of such agents, since it deals only with rational agents who obey the standard probability axioms.

Even when probabilism is seemingly able to account for the behaviour of agents, it does not always do a very good job of it. Consider a rational agent who is certain that r and s are logically incompatible, and on the same stake, the amount of money she is willing to bet on the disjunction $r \vee s$ is the sum of the amount of money she is willing to bet on each disjunct individually. *One* explanation for the agent's betting dispositions is as follows: because she is certain that r and s are logically incompatible, her credence in $r \vee s$ equals the sum of her individual credences in r and in s . For the probabilist, however, there is a more fundamental explanation: because r and s are in fact logically incompatible, and the agent is rational, her credence in $r \vee s$ equals the sum of her individual credences in r and in s . Now this is just Finite Additivity at work, but intuitively, the first explanation seems to be the more fundamental one.

3.2.5 Does not Respect the Implication/ Inference Distinction

Last but not least, probabilism is at risk of ignoring Gilbert Harman's (1986) admonition to respect the distinction between implication and inference. According to Harman (1986), there is a gap between *logic* and principles of reasoning: even if p is logically *implied* by one's beliefs, it does not immediately follow that one ought to *infer* that p on pain of irrationality (pp. 5-6; p. 12; p. 16). Analogously, even if p entails q , for example, it does not immediately follow that one's credence in q ought to be greater than or equal to one's credence in p on pain of irrationality. Harman (1986) does admit, however, that one has reason to believe that p if one recognises that p is implied by one's beliefs; I claim that if one is certain that p entails q , then it is irrational

for one to have a lower credence in q than in p (pp. 17-18).

3.3 Proposed Constraints on Rational Credences

In this section, I shall propose a number of constraints on rational credences that are meant to be an improvement upon the ones imposed by weak probabilism: if one violates them, then one is irrational in an epistemically blameworthy sense. They are also meant to be more realistic than the ones imposed by strong probabilism: they allow agents to fall short of logical omniscience, and more generally, to assign probabilities of less than 1 to the tautologies of the language over which such probabilities are defined.

Weak global probabilism, by claiming that ideally rational agents obey the probability axioms, assumes that such agents are logically omniscient. Although the assumption is dubious, there is a neighbouring claim that is more plausible: namely, given that an agent is logically omniscient, it seems irrational for her to violate Finite Additivity, and some of the other theorems derivable from the standard probability axioms. But where lies the source of this intuition?

Logically omniscient agents have full credence in all logical truths. But one's probabilistic judgements are in the first instance governed by one's judgements of truth values rather than by logic *per se*. For example, Finite Additivity serves only *indirectly* as a rationality constraint on the credences of a logically omniscient agent. If she violates such a constraint, her failure of rationality does not ultimately depend on whether p and q are in fact logically incompatible. Rather it lies in being absolutely certain that p and q are logically incompatible, and yet having a credence in $p \vee q$ that is different from the sum of the credences she has in p and in q individually. Similarly, a person who lacks logical omniscience is no less irrational if she is mistakenly certain

that p and q are logically incompatible, but her credence in $p \vee q$ differs from the sum of her individual credences in p and in q . In fact, you might well think that the person is doubly irrational. For first, she errs in logic, and then she errs in probability.

With the above in mind, consider, as a first step, the following as constraints on our credences:

For any \mathcal{L} , for any p and q in such an \mathcal{L} , and for any P from such an \mathcal{L} to the interval $[0,1]$,

1. $0 \leq P(p) \leq 1$
2. $P(p|q) = \frac{P(p \wedge q)}{P(q)}$ if $P(q) \neq 0$.
3. If $P(p \text{ and } q \text{ are logically incompatible}) = 1$, then $P(p \vee q) = P(p) + P(q)$.
4. If $P(p \text{ entails } q) = 1$, then $P(q) \geq P(p)$.
5. If $P(p \text{ and } q \text{ are logically equivalent}) = 1$, then $P(p) = P(q)$.

Let the sentences in \mathcal{L} represent propositions that an agent is able to grasp. I take it that Constraint 1 holds as a matter of convention. Constraint 2 is just CP from before.¹¹ Normalisation has been discarded, and Finite Additivity, Equivalence, and Entailment have all been replaced. If we take certainty that p and q are logically equivalent to be certainty that p entails q and that q entails p , then strictly speaking, constraint 5 is redundant.¹²

The constraints are similar in spirit to those provided by Hacking (1967).

Hacking's analogue of constraint 3 can be stated as 'If one *knows* that p and

¹¹I have reservations about CP as the definition of conditional probability, for some of the reasons Hájek (2003) offers. However, it can be treated as a constraint on credences even if not as a definition of conditional probability. In any case, CP does not in itself generate any of the problems I have raised against probabilism.

¹²For my purposes, I assume that a credence of 1 amounts to certainty. But although one may have a credence of 1 that a infinitely fine dart thrown at a dart board will not land on some particular point, one need not be certain that it will not land on that point. See also Maher (1993), pp. 135-136.

q are logically incompatible, then $P(p \vee q) = P(p) + P(q)$ ' (p. 320). I.J. Good (1986) has also suggested a constraint similar to constraint 5: 'If you have proved that p is logically equivalent to q then $P(p|r) = P(q|r)$ ', for any r (p. 8). There are several reasons for preferring constraint 3 to Hacking's analogue of it. First, the latter, but not the former, is silent in the case when p and q are in fact logically compatible, but someone is mistakenly certain, and hence does not know, that they are not. Second, Hacking's constraints appeal to the notion of knowledge, but insofar as we are concerned with that aspect of rationality to do with how credences fit together, we do best by making no such appeal. In fact, a Bayesian epistemologist such as Jeffrey (1992) will prefer not to invoke the notion at all. After all, he thinks that traditional epistemology is outmoded epistemology, and that the notion of knowledge ought to be superseded by that of credence. Third, it is at least arguable that one need not be certain that p and q are logically incompatible to know that they are. But if one is not certain that p and q are logically incompatible, then one's credence in $p \vee q$ might well differ from the sum of one's credence in p and in q . Similar reasons count in favour of constraint 5 against Good's analogue of it. If one proves that p is logically equivalent to q , then p is in fact logically equivalent to q . So Good's constraint, unlike constraint 5, is silent in the case when one is mistakenly certain that p is logically equivalent to q but assigns p and q different credences. Also, it is at least arguable that one may prove that p and q are logically equivalent without being absolutely certain that they are. But in such a case, one may well be rational in assigning p and q slightly different credences.¹³

Constraints 3 to 5 are good as far as they go, but they do not go far enough. For example, whenever one is certain that $p \wedge q$ is false (and not also certain that it is true), whether as a matter of logic, metaphysics,

¹³See section 5.4.5 of chapter 5.

happenstance, or what have you, one's credence in $p \vee q$ ought to equal the sum of one's individual credences in p and in q . Also, as Harman (1986) notes, some might doubt that certain agents are sophisticated enough to possess the concepts of logical equivalence, logical incompatibility, or entailment (pp. 18-19). All of this suggests that the following constraints are more fundamental than constraints 3 to 5:

6. If $P(\sim(p \wedge q)) = 1$, then $P(p \vee q) = P(p) + P(q)$.

7. If $P(p \rightarrow q) = 1$, then $P(q) \geq P(p)$.

8. If $P(p \leftrightarrow q) = 1$, then $P(p) = P(q)$.

In paraconsistent logic, both p and $\neg p$ can be true at once, but if we take $\sim p$ to be true just in case p is false and not also true, then even in paraconsistent logic, p and $\sim p$ can't both be true. The antecedent of constraint 6 reads ' $P(\sim(p \wedge q)) = 1$ ' rather than ' $P(\neg(p \wedge q)) = 1$ ' to avoid deeming as irrational someone who, for some p and some q , is certain that $\neg(p \wedge q)$ even while certain that p , that q and that $p \wedge q$. If the antecedent of constraint 6 had read ' $P(\neg(p \wedge q)) = 1$ ', we would have got the absurd result that for such a person to be rational, $P(p \vee q) = 2$. The antecedent of constraint 7 reads ' $P(p \rightarrow q) = 1$ ' rather than ' $P(p \supset q) = 1$ ' to avoid deeming as irrational someone who, for some p and some q , is uncertain that q , but has full credence not only in p , but in $\neg p$, and hence in $p \supset q$.¹⁴ For a similar reason, the antecedent of constraint 8 reads ' $P(p \leftrightarrow q) = 1$ ' rather than ' $P(p \equiv q)$ '. Strictly speaking, constraint 8 is redundant if we take it that an agent is certain that $p \leftrightarrow q$ just in case she is both certain that $p \rightarrow q$ and that $q \rightarrow p$.

Constraints 3, 4 and 5 can be seen as special cases of constraints 6, 7, and 8 respectively, provided we make the following plausible assumptions. To

¹⁴The argument form ' $p, p \supset q$; therefore q ' is not valid in paraconsistent logic.

be certain that p and q are logically incompatible, one has to be certain that $\sim p \wedge q$. To be certain that p entails q , one has to be certain that $p \rightarrow q$. And to be certain that p and q are logically equivalent, one has to be certain that $p \leftrightarrow q$. Hence, it is irrational to violate constraints 3, 4, and 5, *because* it is irrational to violate constraints 6, 7, and 8 respectively. For example, suppose one assigns p and q different credences, but is certain that p and q are logically equivalent. Then one is irrational, but such irrationality has less to do with being certain that p and q are *logically* equivalent than with being certain that $p \leftrightarrow q$.

Constraints 6 to 8 are similar to constraints proposed by Gaifman (2004). Gaifman proposes ‘to represent the agent’s perspective by some finite, large enough set, \mathbf{P} , of sentences’ (p. 103). He assumes, for simplicity, that modus ponens is the single inference rule for the agent, and takes \mathbf{A} , a subset of \mathbf{P} , to contain logical axioms and what he calls ‘subject-specific axioms’ (*ibid.*, p. 104). $P()$ is a function from \mathbf{P} to the interval $[0,1]$, satisfying the following constraints:

G1. $P(p) = 1$, for every p in \mathbf{A} , which is in \mathbf{P} .

G2. If p is in \mathbf{P} , so is $\neg p$, and $P(\neg p) = 1 - P(p)$

G3. If $P(p \wedge q) = 0$, and p, q and $p \vee q$ are in \mathbf{P} , then $P(p \vee q) = P(p) + P(q)$.

G4. If $P(p \supset q) = 1$, and p and q are in \mathbf{P} , then $P(q) \geq P(p)$.

A local proof is a proof from \mathbf{A} that consists of members of \mathbf{P} . For any p , if p is locally provable (if there is a local proof of p), then by G1 and G4, and by modus ponens being the sole inference rule, $P(p) = 1$. By G4 again, we get the following constraint:

G5. For any p and q in \mathbf{P} , if $p \supset q$ is locally provable, then $P(q) \geq P(p)$.

We can see how the constraints are meant to tame the problem of logical omniscience. As Gaifman points out, the fact that $p \supset q$ and $q \supset r$ are locally provable by an agent does not mean that $p \supset r$ is also locally provable by her. For the proof might require the use of a sentence that is not in \mathbf{P} , say the sentence $((p \supset q) \wedge (q \supset r)) \supset (p \supset r)$. But the move requires Gaifman to say that $P()$ is undefined for such a sentence, and I have my qualms about this. Intuitively, an agent may not employ the sentence as an axiom to prove that $p \supset r$ so long as she lacks absolute certainty that it is true—it need not be that she is unable to grasp the sentence or has no opinion about it whatsoever.

Also, for Gaifman, the credence function of an agent who obeys his constraints is a (partial) probability function, whereas this is not necessarily true of an agent who obeys constraints 6 to 8. Some might see this as an advantage of Gaifman's constraints compared to constraints 6 to 8, since his constraints have significant mathematical structure. But as we shall see towards the end of this chapter, this is not necessarily an advantage. Also, I prefer not to include constraint G2, since it deems as irrational someone who assigns both p and $\neg p$ a high credence, or whose credences in p and in $\neg p$ sum up to less than 1. But such a person is not necessarily irrational, since paraconsistent logic allows both p and $\neg p$ to be true, and intuitionistic logic allows one's credences in p and in $\neg p$ to sum up to less than 1.¹⁵

James Hawthorne (2009) has also offered constraints on credences more realistic than those imposed by probabilism (p. 54). In particular, he has offered a more realistic version of the axioms of qualitative or comparative probability. For example, one of the axioms can be read as 'If S is (warranted in being) certain that r and s are materially equivalent, and is (warranted in being) at least as confident that p is true as she is that r is true, then she is (warranted in being) at least as confident that p is true as she is that s is

¹⁵Cf. Priest (2006), pp. 107-109, and Weatherson (2003).

true.¹⁶ Hawthorne, however, retains the axiom according to which if t is a logical truth, then S is certain that t (*ibid.*). He admits that an agent who violates the axiom need not be irrational, but thinks that the axiom can still be held as a normative ideal (*ibid.*, p. 54; p. 57). As I mentioned earlier, I have no quarrel with this, but I am looking for constraints on credences that agents who are not logically omniscient can reasonably be expected to obey.

I have explained why for my purposes, advancing constraints 6 to 8, together with giving up Normalisation, is better than adopting the set of constraints put forward by Gaifman and by Hawthorne. But let us probe more deeply. Since the constraints proposed do not include Normalisation, they can be satisfied trivially by an agent who does not assign a credence of 1 to any proposition whatsoever. However, even such an agent seems irrational if, for example, her conditional credence in p given $p \leftrightarrow q$ differs from her conditional credence in q given $p \leftrightarrow q$. This suggests the following constraints on one's *conditional* credences:

$$9. P(p|p) = 1$$

$$10. P(p \vee q | \sim(p \wedge q)) = P(p | \sim(p \wedge q)) + P(q | \sim(p \wedge q))$$

$$11. P(p | p \leftrightarrow q) = P(q | p \leftrightarrow q)$$

$$12. P(q | p \rightarrow q) \geq P(p | p \rightarrow q)$$

Suppose we add the following plausible-looking constraint:

$$13. \text{ If } P(q) = 1, \text{ then } P(p|q) = P(p).^{17}$$

¹⁶I think that offering constraints on qualitative credences is indeed a worthwhile endeavour, although for the purposes of this dissertation, I assume that credences are always quantitative.

¹⁷This constraint can be obtained from constraint 2 ($P(p|q) = \frac{P(p \wedge q)}{P(q)}$ if $P(q) \neq 0$), with the additional assumption that if $P(p \wedge q) = 1$, then $P(p) = 1$, for all p and q .

We can then derive constraints 6 to 8 from constraints 10 to 13. To derive constraint 6, assume that $P(\sim(p \wedge q)) = 1$. By constraints 10 and 13, $P(p \vee q) = P(p) + P(q)$. Discharging the assumption yields: if $P(\sim(p \wedge q)) = 1$, then $P(p \vee q) = P(p) + P(q)$. To derive constraint 7, assume that $P(p \leftrightarrow q) = 1$. By constraints 11 and 13, $P(p) = P(q)$. Discharging the assumption yields: if $P(p \leftrightarrow q) = 1$, then $P(p) = P(q)$. To derive constraint 8, assume that $P(p \rightarrow q) = 1$. By constraints 12 and 13, $P(q) \geq P(p)$. Discharging the assumption yields: if $P(p \rightarrow q) = 1$, then $P(q) \geq P(p)$.

Constraints 10 to 12, unlike constraints 6 to 8, are not satisfied trivially by agents who do not assign a credence of 1 to any proposition whatsoever. For instance, someone who is not certain that $p \rightarrow q$ vacuously satisfies constraint 7, but might still violate constraint 11. Intuitively, constraints 6 to 13 are also more acceptable than those provided by the probabilist. What justification is there for accepting them? Other than by appealing to intuition, I have no direct justification to offer. But there is plenty of indirect justification. First, note that constraints 6 to 13 can all be derived from Non-negativity, Normalisation, Finite Additivity, and CP. Hence, the probabilist should accept them by her own lights, even though the strong probabilist may question whether the constraints on their own serve adequately as rationality constraints on credences (I address this issue below), whereas the weak probabilist can only say that they are constraints on ideal rationality.

Furthermore, the proposed constraints avoid the problems that plague probabilism. First, since they do not include Normalisation, an agent who assigns a credence of less than 1 to a long and complex logical truth need not be deemed irrational. The constraints thus take the enterprise of logical inquiry seriously, since they allow a rational agent to be uncertain of both logical truths and non-logical truths. Second, for the same reason, the constraints are not too beholden to logic: they are neutral on the issue of which of the

main logics, if any, is the correct one. Third, the constraints provide a better account of how our credences fit together than the probability axioms do. They do not blur the distinction between irrationality due to errors in logic, and irrationality due to errors in probability. They also respect the intuition that one's probabilistic judgements are in the first instance governed by one's judgement of truth values rather than by logic *per se*. Fourth, they allow us to explain the behaviour of an agent better than probabilism explains it. For example, they allow us to explain why an agent who has a credence of 0.8 in t , where t is a logical truth, may be willing to bet against the truth of t in some situations. Also, according to the proposed constraints, when an agent acts in a way that accords with the axioms, for example, with Finite Additivity, the more fundamental explanation for the agent's behaviour is that she is certain that two particular propositions are logically incompatible, rather than that the two propositions are in fact logically compatible. The probabilist, however, is committed to the implausible view that the latter explanation is the more fundamental one. Fifth, the proposed constraints, unlike the constraints imposed by the probabilist, are less likely to violate the implication/ inference distinction.

3.4 Objections and Replies

The advantages listed above may not be enough to convince one to give up probabilism, and adopt the proposed constraints as more fundamental constraints on rational credences than those provided by the probabilist. For one may think that the proposed constraints, taken on their own, face serious problems. I shall now anticipate and answer some objections.

3.4.1 Irrationality Twice Over

First, one might worry that my constraints allow for irrationality twice over. For example, it might already be bad enough for one to be less than certain that two claims that are obviously incompatible are incompatible—but my constraints do not penalise such a person if her credence in the disjunction of the two claims differs from the sum of her credence in each claim.

In response, the constraints are meant to deal with how credences fit together. Admittedly, it is irrational for one to be less than certain that two claims that are obviously incompatible are incompatible. But given that one is already in such a doxastic state, it is not a failure of one's credences to fit together, and one does not suffer greater irrationality, if one's credence in the disjunction of the two claims differs from the sum of one's credence in each claim. To help drive home the point, consider the following analogy. Suppose an exam question has two parts, and the answer to the second part has to be derived from the answer to the first. Suppose also that a student gives the wrong answer to the first part of the question, and as a result, gives the wrong answer to the second part as well. There is no need to penalise her twice over if she has employed the right method in deriving her second answer from the first.¹⁸

3.4.2 Too Undemanding

Second, one might object that the proposed constraints are too undemanding, in the sense that they do not penalise those who assign a credence of less than 1 to *obvious* logical truths. But what counts as an obvious logical truth according to one logic may not be counted as a logical truth according to another logic. Since it is not obvious which logic is the correct logic (if any), what counts as an obvious logical truth according to one logic may not in fact

¹⁸Thanks to John Cusbert for the analogy.

be *obviously* true, even if it is *logically* true.

Perhaps the objection may be modified as follows: in classical logic, some claims are obviously tautologous, and those who are certain that classical logic is correct are downright irrational if they assign a credence of less than 1 to any such claim (setting aside the issue of whether they are rational to be certain that classical logic is the correct logic). My constraints do not penalise such agents.

Also, we earlier encountered a related objection to constraints 6 to 8—one may point out that an agent can satisfy them trivially by not being certain of any claim of the form $\sim(p \wedge q)$, $p \rightarrow q$, or $p \leftrightarrow q$. The introduction of constraints 9 to 12 is meant to reduce the force of such an objection, but one might still think that the constraints I have offered on *unconditional* credences are too undemanding. In fact, one might think that constraints 6 to 8 encourage agents to refrain from being certain of anything whatsoever.

In reply to the first objection, I grant that such agents may well be irrational. But this is not a problem for my proposed constraints. After all, satisfying the constraints is necessary but not sufficient for rationality—probabilists themselves typically take the axioms to offer necessary but not sufficient constraints on rational credences. For instance, someone who assigns a credence of 1 to ‘It is raining but I don’t believe it’ or to ‘I don’t exist right now’ is irrational, even though she may satisfy the proposed constraints, or for that matter, the probability axioms. Besides, there is another way to explain such irrationality, if, as I have urged, we take seriously the enterprise of logical inquiry. As Hacking (1967) suggests, rationality demands that we search for truths, so long as it is cost-effective to do so (pp. 322-323).¹⁹ This helps explain why it is irrational to have a credence of less than 1 in obvious truths, whether logical or non-logical, since being obvious, the search for such truths

¹⁹Hacking attributes the idea to I.J. Good.

is virtually cost-free.

Observe also that if the objection is a strike against my proposed constraints, it's also a strike against weak probabilism. For by itself, all that weak probabilism entitles us to say of someone who does not assign a credence of 1 to an obvious logical truth is that she is not ideally rational. And to explain why such a person is irrational in the epistemically blameworthy sense, the weak probabilist might well appeal to an explanation like the one Hacking suggests.

In reply to the second objection, note again that the proposed constraints are meant to be necessary but not sufficient for rationality. As William James (1956) writes, '*We must know the truth; and we must avoid error,—* [...] they are not two ways of stating an identical commandment, they are two separate laws' (p. 17; James's emphasis). Analogously, it's one thing to strive to increase one's confidence in truths, and another thing to strive for coherence (understood in the intuitive sense rather than in the probabilist's technical sense). The constraints are concerned with the latter; other principles of rationality deal with the former.

3.4.3 Lack of Mathematical Structure

Third, one might object that my proposed constraints cannot do the work that probabilism does. As seen in the previous chapter, Pollock thinks if we impose less demanding constraints on credences than those imposed by probabilism, we lose certain theorems in Bayesian epistemology. Similarly, Talbott (2008) alleges that if we allow for constraints on rational credences that do not demand logical omniscience, we 'block the derivation of almost all of the important results in Bayesian epistemology'. And Garber (1983) writes that '[e]xplicitly relativizing coherence to an individual's state of knowledge [or in my case, certainties or suppositions] might seem attractive [but] will give us little of the mathematical structure we want' (p. 109).

Now I grant that my proposed constraints, as they stand, have significantly less mathematical structure than the constraints imposed by probabilism. But we can easily get the desired mathematical structure out of them by making certain idealisations. Suppose we want to study a particular decision-theoretic situation, and we assume, as most decision theorists do, that classical logic is the correct logic. Suppose also that for simplicity, we assume that the agents involved in such situations are logically omniscient, at least with respect to classical logic. My proposed constraints allow us to make such idealisations to gain mathematical power. To see this, take the following constraint: For any \mathcal{L} , for any p and q in such an \mathcal{L} , and for any P from such an \mathcal{L} to the interval $[0,1]$,

$$6. \text{ If } P(\neg(p \wedge q)) = 1, \text{ then } P(p \vee q) = P(p) + P(q).$$

Now consider Xuxa—let \mathcal{L}_X contain the propositions over which her credences are defined, and let P_X represent her credence function. Suppose that Xuxa is certain, for all p and q in \mathcal{L}_X , whether p is a classical tautology, and whether p and q are logically incompatible according to classical logic. I shall show that P_X is also a probability function, provided that Xuxa does not violate constraint 6. From constraint 1,

$$\text{X1. } 0 \leq P_X(p)$$

By hypothesis,

$$\text{X2. If } A \text{ is classically tautologous, then } P_X(p) = 1.$$

Now constraint 6 dictates that Xuxa obey X3 below:

$$\text{X3. If } p \text{ and } q \text{ are incompatible according to classical logic, then } P_X(p \vee q) = P_X(p) + P_X(q).$$

For suppose that according to classical logic, p and q are logically incompatible, for some p and some q . Xuxa violates X3 if $P_X(p \vee q) \neq P_X(p) + P_X(q)$. But by hypothesis, Xuxa assigns a credence of 1 to $\neg(p \wedge q)$. Hence, a violation of X3 results in a violation of constraint 6 as well. Based on X1 to X3, we can derive various familiar theorems of probability. Given that Xuxa is stipulated to obey X1 and X2, the only way that she will violate these theorems is if she violates X3. But since she has to obey X3 on pain of violating constraint 6, Xuxa has to obey all the derived theorems on pain of irrationality: her rational credence function is a probability function. However, although she is irrational in violating either X3 or the derived theorems, such irrationality is not due to a violation of X3 *per se*; after all, it may turn out that classical logic is not the correct logic. Rather, her irrationality lies in the violation of constraint 6. Here is a further illustration of the point. Since by Xuxa's lights, $P_X(p \vee \neg p) = 1$, and p and $\neg p$ are logically incompatible, she should obey the following theorem: $P_X(p) = 1 - P_X(\neg p)$. But consider an agent who thinks that intuitionistic logic is the correct logic. So she may not be certain that $p \vee \neg p$ is true. So even if by her lights, p and $\neg p$ are logically incompatible, she is not subject to a similar constraint. (Cf. Weatherson (2003).) It need not be irrational for her to violate the theorem because doing so might not lead her to violate constraint 6.

Let's consider one more example to illustrate the flexibility of the proposed constraints. Xuxa is logically omniscient, at least with respect to classical logic. But we can also make slightly weaker idealisations. Consider Yoko, who is cognitively limited, and let Q be a restricted subset of sentences in \mathcal{L}_Y that represent the propositions over which her credences are defined. Suppose that Yoko is certain, for all p and q in Q , whether p is a classical tautology, and whether p and q are classically, logically incompatible. From constraint 1,

Y1. $0 \leq P_Y(p)$.

By hypothesis,

Y2. If p is in Q and is classically tautologous, then $P_Y(p) = 1$.

On pain of violating constraint 6, Yoko should satisfy the following:

Y3. If p , q , $p \vee q$, and $\neg(p \wedge q)$ are in Q , and p and q are incompatible according to classical logic, then $P_Y(p \vee q) = P_Y(p) + P_Y(q)$. (Cf. Gaifman (2004).)

For suppose that the antecedent of Y3 obtains. Then $P_Y(\neg(p \wedge q)) = 1$. If Y3 is violated, then $P_Y(p \vee q) \neq P_Y(p) + P_Y(q)$, which also means that constraint 6 is violated. Note that some other agent even more cognitively limited than Yoko may not be irrational if for some p and q in Q , she assigns $p \vee q$ a different probability from the sum of the individual probabilities assigned to p and to q , provided that in doing so, she does not violate constraint 6.

Now one might want to hold that X3 is a fundamental constraint on Xuxa's credences, and Y3 a fundamental constraint on Yoko's, and there is nothing deeper which explains why Xuxa's violation of X3 and Yoko's violation of Y3 are irrational. But this would not allow us to explain why there seems to be something common to Xuxa's violation of X3 and Yoko's violation of Y3. By holding that their violations are fundamentally violations of constraint 6, we are able to provide such an explanation. The proposed constraints can thus serve as the foundation upon which certain idealised models of rational credences rest. The above examples show that although the proposed constraints lack structure as they stand, this is not a problem, but an advantage. We can get varying amounts of structure out of the proposed constraints by making different idealising assumptions about the kind of agent we are interested in studying.

All this bolsters the claim that we can reject probabilism without losing important results in Bayesian epistemology, *pace* Pollock (2006), Talbott (2008) and Garber (1983). What has led to such important results in the

first place is not the *truth* of probabilism, but the *idealising* assumption that probabilism is true, and we should not conflate the two. We know from science that idealisations can be very useful. No glass surface is frictionless, but when studying the movement of balls rolling down a glass pane, it is sometimes useful to make the idealising assumption that the pane is frictionless. Analogously, when studying real agents, it may sometimes be useful to make certain idealisations. First, note that we can make idealisations that do not demand full blown logical omniscience, but that give us significant mathematical structure—witness the constraints on rational credences offered by Garber (1983) and Gaifman (2004). Second, note that even the assumption of full blown logical omniscience may be useful when we are dealing with cases in which an agent’s lack of logical omniscience is unlikely to manifest itself. When trying to predict and explain the actions of scientists and people making deliberate decisions, it is sometimes reasonable to operate on the guiding assumption that it is not too hard for them to figure out the truth or falsity of the logical claims they have to work with, given sufficient resources. In such cases, it is also safe to assume that after deliberation such people will do just what an idealised agent who is logically omniscient will do. Of course, when we have reason to suspect that such assumptions are unwarranted, we should not use the probability axioms blindly to predict how such people will act. Instead, we can turn to constraints on credences that are less demanding than those provided by probabilism.

Also, even though an agent whose credences violate the standard probability axioms need not be epistemically blameworthy, there is still a good sense in which the axioms serve as a guide to rational credences. Consider the case of how logic is often taken to serve as a guide to rational binary beliefs. We can accept Harman’s implication/ inference distinction and grant that it might not be irrational for an agent to have two beliefs that are inconsistent. But

suppose she notices the inconsistency, and consequently, gives up one or more of the beliefs. If asked to justify her change in mind, she may offer the justification that the beliefs are inconsistent—and not merely that she believes that the two beliefs are inconsistent. Analogously, suppose an agent has different credences in p and in q initially, comes to be certain that p and q are logically equivalent, and consequently assigns equal credences to p and q . If asked to justify her change in mind, she may offer the justification that p and q are logically equivalent—and not merely that she is certain that they are logically equivalent. So the probability axioms can be seen to provide a probability logic that we can draw upon to help justify our having certain credences, even though violations of such a logic need not mean that one is irrational in any epistemically blameworthy way.²⁰ And such a probability logic may be enough for us to retain the important results in Bayesian epistemology and confirmation theory that Pollock and Talbott think we lose if we reject probabilism. After all, analogously, if logic delivers the result that some very complicated argument has a deductively valid form, we retain the result even if we maintain that one need not be irrational in believing (in the binary sense) all the premises of the argument but not its conclusion.

²⁰See Adam (1998) for more on probability logic. Ramsey (1926) is explicit that in arguing for a logic of consistency for graded belief, he is not saying that such a logic coincides with ‘human logic [...], which tells men how they should think, [and which] is not merely independent of but sometimes actually incompatible with formal logic’ (p. 78; 87).

Chapter 4

Credences Do All the Heavy Lifting

4.1 The Sense in which Credences Do the Heavy Lifting

As seen in chapter 2, binary beliefs cannot do all the work of credences when it comes to the explanation and prediction of action and behaviour. Since credences are more fine-grained than beliefs, it's also natural to think that they do all the work that binary beliefs may have been thought to do. In slogan form, credences do all the heavy lifting. I shall clarify what I mean by the slogan in what follows. I shall then argue against some views according to which beliefs do work that credences don't do.

The slogan is concerned with the explanation and prediction of action and behaviour *at the intentional level*. Credences earn their keep at such a level by playing a certain role: in combination with one's desires and other mental states, one's credences determine how one is disposed to behave, and what mental states one is disposed to enter into. For example, suppose that I've a high credence that uncaged tigers are dangerous. This, in conjunction with a desire to avoid danger, a high credence that the best way to do so is to run away, and other mental states such as a feeling of fear, will tend to result in running away behaviour among other things.

What is the place of binary beliefs in the picture above? If all the explanatory and predictive work at the intentional level can be done by credences

and desires (perhaps plus other non-doxastic states), then binary beliefs need not enter the picture at all, at least not as something over and above credences and desires. One might think that binary beliefs enter the picture via playing a role that supervenes on, among other things, the roles that credences and desires play. But even if so, their place will be peripheral—any role that they play will be explanatorily redundant, because whenever we appeal to it, we can instead appeal to the roles played by credences and desires.

Frankish (2009) claims just as an intentional explanation of an action is not made redundant by the availability of a physical explanation of it, an intentional explanation of an action in terms of binary beliefs is not made redundant by the availability of an intentional explanation of it in terms of credences. (pp. 88-89). By holding that beliefs are explanatorily redundant at the intentional level, am I not proving too much? Am I committing myself to the claim that only physical terms do any real work in the explanation and prediction of action and behaviour, given that credences supervene upon neurophysiological properties that in turn supervene upon microphysical properties?

I'm willing to grant (but refrain from endorsing) the claim that all things considered, only microphysical properties, perhaps in conjunction with natural laws, do any real work. But credences and microphysical properties are invoked in *different* kinds of explanation, whereas binary beliefs and credences are both invoked in the *same* kind of explanation. My claim is that insofar as we're interested in an *intentional* account of action and behaviour, credences rather than beliefs do all the work.

At the intentional level, either binary beliefs and credences compete to play certain roles in explaining and predicting action and behaviour, or they offer us two compatible ways—one coarse-grained and one more fine-grained—to cash out these roles. If the first, then some might think that the share of the roles is divided between credences and binary beliefs. But in fact,

credences will muscle out binary beliefs, since there is work that credences do that binary beliefs cannot do, as argued in chapter 2, and any work that binary beliefs purport to do can be done by credences, as shall be argued below. If the second, then there is no non-redundant intentional role that binary beliefs play. At the intentional level, whenever we provide an account of action and behaviour in a coarse-grained way, we can also do it in a fine-grained way, but we can't always replace a fine-grained account with a more coarse-grained one.

I do not hold that we never appeal to or that it is never useful to appeal to binary beliefs when explaining or predicting action and behaviour. But such appeals are typically quick and dirty. Sometimes the quick and dirty are all we need; perhaps for lack of time or patience, we might not even want explanations and predictions that are too detailed or too accurate. But this just shows that appealing to the notion of binary belief is sometimes *expedient*.¹

One might point out that an explanation in coarse-grained terms may sometimes be the only proper explanation available. As Putnam (1973) points out, in explaining why a square peg can pass through a square hole but not a round hole on the same board, it is more illuminating to appeal to the macrophysical properties of the board and the peg than to appeal to their microphysical properties (p. 132). Pointing out that the peg and the board are both rigid, and that the square hole but not the round hole is big enough for the peg to pass through, yields a perfectly good explanation. But 'the deduction

¹Cf. Lewis (1996) on knowledge ascriptions; he writes:

If you doubt that the word 'know' bears any real load in science or in metaphysics, I partly agree. The serious business of science has to do not with knowledge *per se*; but rather, with the elimination of possibilities through the evidence of perception, memory, etc., and with the changes that one's belief system would (or might or should) undergo under the impact of such eliminations. Ascriptions of knowledge to yourself or others are a very sloppy way of conveying very incomplete information about the elimination of possibilities. [...] The only excuse for giving information about what really matters in such a sloppy way is that at least it is easy and quick! But it *is* easy and quick; whereas giving full and precise information about which possibilities have been eliminated seems to be extremely difficult [...]. (p. 563)

of the fact to be explained from the positions and velocities of the elementary particles, their electrical attractions and repulsions, etc.—fails to *explain*' (*ibid.*; Putnam's emphasis). Too much irrelevant information is involved.

Analogously, an explanation of one's behaviour in terms of binary beliefs rather than credences might provide just the right amount of detail.² Why is Jack carrying an umbrella? Given that he has a great desire not to get wet, and that he is very confident that carrying an umbrella when it's raining will help him satisfy this desire, saying that he believes that it's raining might be a perfectly fine answer; saying that his credence that it's raining is 0.95 might say too little by saying too much—he would presumably have carried his umbrella even if, say, his credence in rain had been 0.97 or 0.98.

I take the point that in explaining someone's action and behaviour, it is not always illuminating to cite the exact credence the person has. But we can avoid being overly precise by talking about credences in coarse-grained terms. For example, we can just say that Jack is very confident that it's raining or that he's more confident that it's raining than that it's not. Or we can report that his credence that it's raining falls roughly in the interval between 0.9 and 0.99 (for example). Or we can say that his expected utility in carrying an umbrella is higher than his expected utility in not carrying an umbrella, without reporting the exact values of those expected utilities. In saying such things, we need not invoke talk of binary belief.

Some philosophers, however, think that binary beliefs have a non-redundant role to play in the explanation and prediction of action and behaviour over and above the roles played by credences and desires. I shall classify such views into three broad categories. According to views in the first category, the very notion of credence is dependent upon the notion of binary

²Thanks to Alan Hájek and Declan Smithies, each of whom pressed this point independently (personal communication).

belief. According to views in the second category, binary beliefs have a non-redundant role to play in reflective, deliberate reasoning. According to views in the third category, binary beliefs have a non-redundant role to play in our practice of assertion.

4.2 The Supposed Dependence of Credence on Binary Beliefs

Views in the first category have been propounded by Isaac Levi (1980), Robert Nozick (1993) and Mark Norris Lance (1995) (although Lance talks about acceptance rather than binary beliefs). If such views are correct, we cannot even appeal to credences to explain and predict our actions and behaviour unless we presuppose that we have certain binary beliefs. I shall go through each view in turn.

4.2.1 Levi's view

According to Levi (1980), to determine our credence in p , we first need to determine the set of *serious possibilities*, which are in turn determined by what one knows or fully believes. Our unconditional credence in p equals our credence in p , conditional upon any proposition entailed by the set of serious possibilities. Since for any proposition e entailed by the set of serious possibilities, our conditional credence in e given e is 1, our unconditional credence in e is 1, for all e (Levi, 1980, p. 76).³

Although Levi is explicit about not trying to explicate the notion of a serious possibility, he offers some examples to illustrate what he means by it (*ibid.*, p. 5). For instance, suppose that at time t , S knows or fully believes that the earth will not explode, or that a coin that is tossed will not fly out

³In slightly greater detail, the set of serious possibilities is represented by K , a consistent set of sentences in the language L (p. 76). One's credences are represented by a function $Q(h; e)$, defined for all h in L , and all e in a subset E of L , where E is a set of hypotheses consistent with K . When K entails e , $Q(h; e) = Q(h)$, where $Q(h)$ is the unconditional probability that h .

into space towards Alpha Centauri (p. 3). Then neither scenario, although logically possible, is a serious possibility. Levi concedes that there might well be times when S takes such possibilities seriously, and assigns positive credence to them. But he holds that in such cases, S is a different cognitive state from his cognitive state at t (*ibid.*).

Why appeal to the notion of serious possibility, and not always let some credence be assigned to logical possibilities, say, to the logical possibility that the earth will explode? Levi thinks that if we don't appeal to serious possibilities, we run the risk of never being certain of anything contingent. But he holds that '[w]hoever denies he is certain of anything is neurotic' (*ibid.*, p. 74). Also, he thinks that knowledge and full beliefs have an important role to play in deliberation and inquiry, and being an infallibilist, he thinks that one has to be certain that p to know that p or to have a full belief that p . But if we are never certain of contingent propositions, then there is nothing contingent that we know or fully believe, and this will in turn diminish the role of knowledge and full beliefs in deliberation and inquiry. (*ibid.*, p. 13).

I am inclined to agree with Levi that anyone who denies that he is certain of anything (contingent or otherwise) is neurotic.⁴ But this is because I'm inclined to think that in ordinary language, to be *certain* that p need not amount to having a credence of 1 in p defined over all the logical possibilities that one is able to carve out. Rather, the word 'certain' is context-sensitive—in some contexts, one needs to have a credence of 1 in p to count as being certain that p , but in other contexts, one doesn't.⁵ A version of the view that the word 'certain' is context-sensitive has been suggested by Lewis (1979), and another version has been defended by Jason Stanley (2008) (pp. 353-354; pp.

⁴Levi's use of the term 'neurotic' is tendentious. As Alan Hájek has pointed out to me, anyone who claims to be certain of anything (contingent) might well be called *dogmatic*!

⁵In various places in this dissertation, I use the word 'certain' to refer to full credence—the context should make it clear when I am doing so!

54-55). If Lewis and Stanley are right, then we may count as being certain that propositions such as ‘The earth will not explode’ or ‘A coin that is tossed will not fly out into space towards Alpha Centauri’ are true, even if we do not assign full credence to them.

In any case, letting credences depend on serious possibilities in the way Levi envisages comes at some cost. Take any mundane contingent proposition, p . On Levi’s view, if $\neg p$ isn’t a serious possibility at time t , and S fully believes that p , then she has full credence in p . But suppose that I’m invariably disposed to reject any bet on the truth of p that pays me one cent if p , but subjects the entire human race to eternal torture if $\neg p$. I do not have full credence in p , which on Levi’s account, means that there is no point in time at which I know or fully believe that p .⁶ But it’s implausible that for me to know or to believe that p , where p is some mundane proposition such as ‘Canberra is the capital of Australia’, I must not be invariably disposed to reject any bet such as the one above.⁷

Levi might reply to the objection by holding that at time t , and before the bet is offered to me, I might indeed be certain that p . However, when it is offered to me, my cognitive state changes. The possibility that $\neg p$ becomes serious, and my credence in p becomes less than 1. Now admittedly, the very offer of a bet might sometimes change one’s cognitive state regarding the truth of p . For example, the offering of a certain bet on p might itself be evidence for or against the truth of p . But we can stipulate that the case above is not like this. (Suppose that before the bet is offered, my credence in p is the same

⁶I am not advocating a betting interpretation of credences, on which to have a certain credence in p just is to be disposed to bet on the truth of p in certain ways. However, as mentioned in chapter 1, I assume that when one has a certain credence in p , one’s credence, together with one’s desires, credences about other propositions, etc., typically disposes one to act in certain ways (that might or might not involve bets).

⁷Someone who subscribes to Bayesian conditioning or to Jeffrey conditioning as the right account of belief revision might worry that neither account will be able to handle an agent’s change in mind that is due to a change in what she takes to be serious possibilities. But Levi (1980) attempts to provide new rules for belief revision.

as my credence in p conditional on the bet being offered to me.) Levi will then have to hold that even though the offer of the bet would not give me any evidence for or against p , it would somehow put me in a cognitive state in which the standard for counting possibilities as serious is different from before. Such a change would presumably explain how I can be certain that p at time t , and yet would reject the bet if it were offered to me.

However, such a reply misses the point of the objection. Entering into a new cognitive state takes time. Admittedly, at a time later than t , I might not be disposed to reject the bet. But the point is that *at time t* , I'm disposed to be such that if I had been offered the bet at that time, I would have rejected it. And this disposition seems to be best explained by holding that I've *some* credence in $\neg p$ *even at t* . Either Levi has to admit that he is not able to capture this aspect of my cognitive state at t , or he has to hold that it is a serious possibility for me that $\neg p$ even at t . But if most people are like me, the latter move is tantamount to holding that there are many mundane contingent propositions in which they do not have full credence after all.

One might think that my disposition at t to reject the bet if it were offered to me is not due to my having some credence in $\neg p$ at that time, but due to my being disposed to acquire a new cognitive state if I should be offered the bet. This is a *possible* explanation, no doubt, but doesn't seem to be the *best* one. Assuming that I'm rational, it's hard to see why I will be disposed to acquire this new cognitive state. After all, it has been stipulated that merely being offered the bet does not lead to any new evidence relevant to the truth of p .

4.2.2 Nozick's view

Nozick (1993), like Levi, thinks that credences are dependent upon what one believes. He doubts that we can formulate a coherent position according to

which we ‘assign probabilities to each and every statement without definitely believing any one, and, in choice situations, act upon these probabilities by (perhaps) maximising expected utility’ (p. 94-95). He writes:

a choice situation is one a person *believes* he is facing, that is, one where he believes that he can perform various alternative actions A_1, \dots, A_n , believes that A_1 can have various possible outcomes O_i (perhaps depending upon which state of the world S_j obtains), and so on. To be sure, the person then acts upon probabilities, $\text{prob}(O_i/A_1)$ or $\text{prob}(S_j/A_1)$, but these occur *within* a structure of beliefs about the situation of choice. [...] Without those beliefs about the situation he was in, his choices would not indicate [his] probabilistic judgements. The theoretical definition of [the latter notion] presupposes attributing certain beliefs to the person. (Nozick, 1993, p. 95)

Setting aside Nozick’s own account of binary belief (of which more later), it is puzzling what an appeal to binary belief is supposed to do. Even if we have a binary belief that we are facing choice situation C_1 , it seems sensible to ask whether the belief is accompanied by a corresponding credence of maximum strength. If so, then appealing to the belief to determine our choice situation is redundant. We can just say that the choice situation is one that we are 100% confident we are facing. But suppose the belief need not be accompanied by full credence. Say, for example, that we believe that we are facing choice situation C_1 with a credence of 0.9, with the rest of our credence invested in the proposition that we are facing choice situation C_2 , the proposition that we are facing choice situation C_3 , and so on, where C_2 , C_3 , and so on are choice situations that contain different sets of possible actions. In such a case, our decision regarding which act to perform does not just depend on the choice

situation we *believe* we are facing. The following explains why.

Note that even Nozick should grant that we might be wrong about the choice situation we are facing. Suppose that we have a binary belief that that our only two options are to walk home or to ride home on our bicycles, and this belief is accompanied by a corresponding credence of 0.9. Suppose also that we have a credence of 0.1 that our only two options are to take a bus home or to take a train home, and the prospects of walking or riding home are extremely dreary, whereas the option of taking a bus or a train home is very welcome. In such a case, we might take a calculated risk and try either to take a bus home or take a train home. If, as luck would have it, these two options are indeed available, then we might well succeed in going home either by bus or by train. But if we only act in a way compatible with the choice situation we *believe* we are facing, then we would not have done so.

One might reply that in the above case, we don't really believe that our only two options are: walk home or ride home. This reply, however, seems to concede that it's credences rather than beliefs that drive the way we act. And for reasons given in chapter 2, it won't help to hold that what we really believe is that there is an objective probability of 0.9 that our only two options are to walk home or to ride home—we cannot in general replace credence talk with talk of binary beliefs about objective probabilities.

Nozick (1993) also thinks that a point similar to the one he makes regarding choice situations can be made about Dutch Book Arguments (used to support the claim that rational credences obey the probability axioms). He claims that for Dutch Book Arguments to work, the agent who figures in them needs to *believe* that she is in such-and-such a betting situation, with such-and-such betting odds, and such-and-such stakes (p. 95). Suppose the agent is in Betting Situation A, but although she has high credence that she is in Betting Situation A, she also has some credence that she is in Betting Situation B.

In such a case, ‘over many such situations she will not conform to the axioms of the probability calculus with respect to the statements that appear (to the outside observer) to be the only relevant ones in [Betting Situation A]’ (*ibid.*, p. 95).

But if someone is unsure whether she is facing Betting Situation A or Betting Situation B, why not just describe her as being in a betting situation in which she is unsure what the odds and stakes are? We can then conduct the Dutch Book Argument relative to this new betting situation. Also it is unclear why appealing to binary beliefs will help someone who wishes to make such arguments. If the binary beliefs that we supposedly have to invoke in Dutch Book Arguments need to be accompanied by full credences, then they are redundant—we can just appeal to full credences. If, however, the relevant beliefs need not be accompanied by full credences, then the supposed problem described in the previous paragraph arises, and the appeal to such beliefs is useless.

Now Nozick (1993) endorses tentatively the view that belief is accompanied by a credence of 1 spread over statements that are not ignored, where which statements are ignored depends on contextual features having to do with what is at stake or with certain intellectual or epistemic standards (pp. 96-97). When we make decisions, and believe that we are facing a particular choice situation, there are various other choice situations that we just ignore or are certain we do not face; when an agent figuring in a Dutch Book Argument believes that she’s in a certain betting situation, she ignores various other betting situations or is certain that she’s not in them. But Nozick’s account faces a problem similar to one that Levi’s account faces. Consider a bet that pays me one cent if I’m facing a certain choice situation, C_1 , but that subjects the entire human race to eternal torture if not. To account for my invariable disposition to reject such a bet, it seems best to hold that in every context, I

have some credence that I am not facing C_1 . But on Nozick's view, this would mean that there is no context in which I believe that I'm facing C_1 , even if C_1 is just some mundane and contingent choice situation.

4.2.3 Lance's view

Lance, unlike Levi and Nozick, talks about acceptance rather than binary belief. He notes that 'a number of philosophers have argued that Bayesian decision theory and epistemology vitiate the status of belief or, in their semi-technical terminology [...], "acceptance"' (p. 147). He then goes on to say that *pace* such philosophers, 'Bayesian decision theory [taken as a prescriptive enterprise] is itself conceptually dependent upon the notion of acceptance'.⁸

On the face of it, it seems that Lance takes acceptance to be somewhat similar to belief. But he does not. According to him,

[Acceptance is] a state of commitment to a proposition which leads, in appropriate circumstance, to acts of asserting which are themselves socially recognised as conferring justificatory commitments—commitments to defend the proposition against legitimate criticism—upon the assertor. (p. 150).

If one can commit oneself to defending a proposition even if one does not believe it, then acceptance is not belief. It is also clear from other textual evidence that Lance does not take acceptance to mean belief. In his paper, he is explicit about not conflating the two (p. 157). He also notes that one could accept a proposition in which one has a credence less than 0.5—but presumably, one does not believe something in which one has such low credence (p. 160).

However, even though Lance is concerned with acceptance, and not with belief

⁸Lance is concerned with prescriptive (or what he calls normative) decision theory, but I take it that if prescriptive decision theory is conceptually dependent upon the notion of acceptance, then the notion is an important one to have even when we are using decision theory to describe how agents who are more or less rational behave.

per se, it'll still be interesting to see whether his argument can be modified so as to apply to binary belief. And there are philosophers like Mark Kaplan (more on him below) who think that one can believe something in which one invests a credence of less than 0.5. Such philosophers might well think that what Lance calls *acceptance* deserves the name *belief*.

Lance (1995) does not think that agents can or ought to carry out expected utility calculations every time they have to make a decision (p. 151). But he thinks that in order for decision theory to be a potentially useful instrument for telling the decision theorist which course of action is rational for an agent to adopt, it should at least be possible for decision theorists, *qua* theorists, to carry out the relevant calculations (*ibid.*). However, he thinks that 'Bayesian decision theory cannot [characterise] rational decision for an agent without presupposing propositions which are accepted by that agent'—without the notion of acceptance, expected utility calculations will be intractable for any theorist who is a finite being (*ibid.*, p. 155).

I shall discuss an example Lance uses to support his claim, but modify it so that it does not assume, as his original example does, that the agent figuring in the example is also the theorist.⁹ Suppose that Jack is deciding whether to go skating, but is concerned about whether the ice is too thin. To figure out what he ought to do, we determine that for him the utility of skating on thick ice is +5, the utility of not skating on thick ice and the utility of not skating on thin ice are both -1, and the utility of skating on thin ice is -20. Lance asks whether our determination of such utilities would be reasonable if there were snipers hiding in the woods intent on shooting any skater who does not fall through the ice (presumably assuming that Jack knows about this).

⁹The modification helps make it clear when the discussion focuses on the point of view of the agent, and when it focuses on that of the theorist. As Lance himself distinguishes between the agent and the theorist, such a modification should neither be biased in favour of him nor against him.

He thinks that the answer is obviously no, and observes that our determination of utilities is defeasible—new information might change what we determine to be the utility of a certain state of affairs. Based on this observation, Lance suggests that we are justified in determining the utility of skating on thick ice to be +5 only if Jack *accepts* that there are no such snipers.

It is natural to reply that Jack need not accept such a thing; he merely needs to have very high credence in it. But Lance thinks if we take the possibility of the presence of snipers into consideration, then to determine what is rational for Jack to do, we need to determine the following utilities: that of skating on thick ice with a sniper present, that of skating on thick ice with no snipers present, that of skating on thin ice with a sniper present, and that of skating on thin ice with no snipers present.

But now, according to Lance, determining these new utilities requires us to suppose that Jack accepts other things. For instance, the utility of skating on thick ice with a sniper present depends on whether the sniper's gun is loaded or not. Hence, determining the new utilities might require us to assume that Jack accepts something like 'The gun of any present sniper will be loaded'. And so on.

Although Lance's argument is meant to show that prescriptive decision theory needs the notion of acceptance, one might think that the argument shows that it needs the notion of binary belief. According to Lance, we may accept some propositions in which we have low credence. But suppose that in the skating example above, Jack accepts the proposition that there are no snipers present even though he has low credence in it. If he acts according to what he accepts, and goes skating, decision theory should deem his action irrational. Hence, one might think that Jack needs to have a binary belief that there are no snipers present for him to be rational to go skating, where to have a binary belief in a proposition requires one to assign it a high enough

credence.

However, Lance's argument for why Jack has to accept some proposition at some point to prevent the calculation of expected utilities from growing too complex is puzzling. When trying to determine whether Jack ought to go skating, given his concerns about the ice being too thin, why not take the utility of skating on thick ice (or thin ice) to be an expected utility that has already factored in Jack's desire not to be shot by a sniper and his low credence that a sniper is present?

The worry is that in order for us theorists to determine the expected utility of skating on thick ice for Jack, we will have to determine the utility of skating on thick ice with a sniper present and the utility of skating on thick ice with no snipers present. And to determine the latter utilities, we will have to determine yet another set of utilities—and so on. But the worry is ill-founded. We might learn that a person is confident that a particular disjunction is true, without learning her opinion about any particular disjunct. We might find out, for example, that she is willing to bet a large sum of money on the truth of the disjunction, without knowing how much she is willing to bet on each individual disjunct. Similarly, we may be able to determine the extent to which Jack is disposed to skate on thick ice, without knowing individually the extent to which he is disposed to skate on thick ice with a sniper present, and the extent to which he is disposed to do so without a sniper present. There is no need for Jack to accept that there are no snipers for us to determine the utility he assigns to skating on thick ice.

Although Lance is right that our determination of utilities is defeasible, the point is a red herring. For example, our determination that someone has low credence in a disjunction is defeasible—we may learn that she has high credence in one of the disjuncts, in which case we should revise our estimation of her credence in the disjunction. It does not follow that in determining her

credence in the disjunction, we need to determine her credence in each disjunct. Similarly, our determination of the utility Jack assigns to skating on thick ice is defeasible—we may learn that Jack is in fact very confident that snipers are present, in which case we should revise our estimation of the utility of skating on thick ice. It does not follow that in determining such a utility, we have to determine individually the utility of skating on thick ice with a sniper present and the utility of skating on thick ice with no sniper present.

4.3 Binary Beliefs and their Putative Role in Reasoning

I have argued that attempts to forge a strong conceptual connection between decision theory on the one hand and the notion of belief (or acceptance) on the other do not work. Let us now look at views that try to carve out a non-redundant role for binary beliefs to play without trying to forge any such conceptual connection.

4.3.1 Practical reasoning

Nozick (1993) has an argument, aside from the two considered in section 4.2.2, for why we need beliefs. He writes:

[t]he task of assigning probabilities to each and every well-formed statement and combination of statements is overwhelming. Beliefs can cut down on this task if probabilities need not be assigned to whatever is incompatible with your beliefs—if all these are automatically ignored or assigned a probability of zero. (Nozick, 1993, p. 96)

Nozick's claim regards the assignment of probabilities to each and every statement as if such an assignment is something performed at a reflective or deliberate level. But recall the distinction made in chapter 1 between decision theory

construed as a way to explain and predict the action and behaviour of agents, and decision theory utilised as a tool by an agent to help her make decisions. The appeal to credences to explain and predict action does not require that people calculate expected utilities deliberately and reflectively before choosing what to do. (It had better not, since we often act without calculating our expected utilities deliberately and reflectively.) This is perfectly compatible with our being disposed to act in ways that maximise or satisfice expected utility. As an analogy, consider vision—massive computations take place in the brain so that we can see, but we do not need to perform such computations consciously.

Perhaps the worry is as follows. Suppose to have a credence in p is to be in a certain mental state. The worry might be that it is psychologically unrealistic for us to act in accordance with decision theory (even an unorthodox decision theory with the more realistic constraints on credences that I proposed in the previous chapter), because there would be too many credences to store and manipulate in our brains.¹⁰

However, recall from chapter 1 that according to the dispositionalist and the interpretationalist, beliefs (whether graded or binary) are not robust mental states that are stored in the brain. So the worry does not arise for them. But how should a functionalist for whom beliefs are robust mental states respond to the worry?

The functionalist may hold that the way beliefs are represented in the brain resemble the way *maps* encode information more than the way that *sentences* encode information.¹¹ A simple map may represent a vast amount of information that will take many sentences (or perhaps a few very long sentences) to represent. For example, consider a typical map representing the

¹⁰Thanks to Declan Smithies for pressing the worry.

¹¹See Braddon-Mitchell and Frank Jackson (2007) for more on the map-like representation of thought versus the sentence-like representation of thought.

locations of various towns and the distances between them. Imagine a line on the map that goes from one spot (representing one town) to another (representing a second town). The various points along this line can be thought of as representing various physical bits of terrain that we have to traverse to go from one town to the other, following the passage indicated by the line. As Braddon-Mitchell and Jackson (2007) point out, there is some truth in the adage that a picture is worth a thousand words.¹² If credences are stored in the brain in a map-like way, the worry that there will be too many credences to store becomes less acute.

Thinking of representation as map-like also enables one to handle the worry that the brain will find it too hard to manipulate the credences stored in the brain. Suppose we represent a new town on our map above. Straightaway, for every other town represented on the map, we get a representation of the distance and direction from that town to the new town. Shift the new town around on the map, and the representation of such distances and directions changes instantaneously. If the brain's manipulation of credences is like manipulating how towns are placed on the map, then just as how shifting towns on the map leads easily to a new representation of the locations of the towns and their respective distances and directions from one another, an increase or decrease in one's credence in a particular proposition may easily lead to changes in the credences and expected utilities that one assigns to other propositions.

But suppose credences are not stored in the brain in a map-like way. The functionalist still has the resources to handle the worry that there are far too many credences for the brain to store and manipulate. Recall the example from chapter 1 in which Ida seems to have the belief that she has no publications, that she has no more than one publication, that she has no more

¹²In fact, if Philip Kitcher and Achille Varzi (2000) are right, some pictures are worth 2^{\aleph_0} sentences.

than two publications, etc. One might worry that there are just too many of such beliefs for them to be stored in Ida's brain. Hence, even a proponent of the view that *binary* beliefs are encoded in the brain might hold that there is a need to distinguish between binary beliefs that are explicit and those that are implicit, where the only the former are encoded in the brain. (In the case of Ida, perhaps only the belief that she has no publications is explicit.) But now, one who is sympathetic to the foregoing should also allow that there is a distinction between explicit and implicit credences, where only the former are stored in the brain. This helps alleviate the worry that there might be more credences than the brain can store.

One might, however, think that only explicit credences are causally efficacious with respect to action and behaviour, and worry that whenever we act, the capacity of our brains is too small to accommodate all such credences. For example, suppose that there is a prize for guessing correctly a positive integer recorded on a piece of paper. Suppose we are told that the number does not exceed z , for some extremely large z . One might worry that for us to come to a decision about our guess, our credences regarding the various contingent propositions, 'The number is 1', 'The number is 2', ..., 'The number is z ', all have to be made explicit. But if z is a very large number, the worry is that there would be too many of such contingent propositions for our credence in each to be stored explicitly in our brains.

But really, we only need to have explicit credences concerning certain equivalence classes of propositions. For example, when trying to decide on our pick, we might just need an explicit credence of 0.7 that the number is anything between 1 and 1 000, and an explicit credence of 0.3 that the number is anything above 1 000, to help us narrow down our pick to any integer in the interval $[1, 1\ 000]$.

So far, I've talked about a kind of decision-making that need not take place at a reflective and deliberate level. But there are some philosophers who hold that binary beliefs have a special role to play in reflective and deliberate reasoning. Timothy Williamson (2009) says the following about binary or what he calls outright belief:

Intuitively, one believes p outright when one is willing to use p as a premise in practical reasoning. Thus one may assign p a high subjective probability without believing p outright, if the corresponding premise in one's practical reasoning is just that p is highly probable on one's evidence, not p itself. (p. 99)

Ralph Wedgwood (2008) has a similar view. He thinks that 'the central difference between a mere partial belief and a full outright belief lies in their different *functional roles*' (p. 4; his emphasis). According to him, this

difference lies in the different kinds of reasoning that these two kinds of belief normally dispose one to perform. If one has a full outright belief in a proposition p , one will simply take p for granted, and treat p as a starting point for further reasoning. [...] By contrast, when one has a mere partial belief in a proposition q , one will not take q for granted in this way [...]. [O]ne will consider, not only arguments that start from q , but also arguments that start from these other incompatible propositions as well; one will in a sense 'weight' each of these arguments by the degree of partial belief that one has in the conjunction of its premises; and one's overall response to these arguments will reflect the weights that one gives them' (*ibid.*)

Continuing the theme, Declan Smithies (forthcoming) writes,

If one is highly confident that p , but one suspends belief that p , then one will be disposed to give some weight in reasoning to the possibility that p is false. [...] By contrast, if one believes that p , then one will be disposed simply to take it for granted that p as a premise in reasoning.

And finally, in a similar vein, Richard Holton (2008) holds that

all-out beliefs enable us to resolve epistemic uncertainty in order to facilitate action. They allow us to reduce an unmanageable amount of information to a manageable amount by excluding certain possibilities from our practical reasoning. (p. 37)

As mentioned, an agent need not calculate her expected utilities deliberately or reflectively before deciding on a course of action. But some philosophers take practical reason as ‘the general human capacity for resolving, through *reflection*, the question of what to do’, and theoretical reason as involving ‘*reflection* with an eye to the truth of propositions’ (R. Jay Wallace, 2009; my emphases). Also, the above quotes from Williamson, Wedgwood, Smithies, and Holton hint at a kind of reasoning that is reflective, deliberate, and mostly conscious. For example, they talk about *using* p as a premise in one’s reasoning, about *considering* arguments that start from certain premises, about *taking* p as a starting point in arguments, and about one’s *weighting* of premises.¹³

Although Williamson, Wedgwood, Smithies, and Holton all want to tie binary beliefs to reasoning, Williamson talks about practical reasoning specifically, whereas the others talk about reasoning in general. In particular, Smithies and Holton are explicit about wanting to tie beliefs not just to practical reasoning but to theoretical reasoning as well. I shall first discuss the

¹³And if they are not talking about deliberate and reflective reasoning, then the points made in response to Nozick at the beginning of this section apply.

view that binary beliefs have a role to play in practical reasoning, and then the view that they have a role to play in theoretical reasoning.

The view under consideration holds that to account for reasoning, we need to appeal to binary beliefs. If this is right, then binary beliefs perform some of the heavy lifting in the explanation and prediction of action and behaviour after all. For reflective reasoning is an act that we perform, and it involves other deliberate acts such as the act of using p as an unqualified premise in such reasoning or the act of computing our expected utilities with pen and paper—we can always choose to employ p with certain qualifications and we can always choose not to compute expected utilities.

But though it is true that we engage in such acts of reasoning, we do not need to appeal to binary beliefs to account for them. Again, think of the distinction between using decision theory as a way to explain and predict action and behaviour and using it as a tool to help us decide what to do. Construed the first way, decision theory does not say that agents always have to calculate expected utilities before acting. However, there are times when we calculate expected utilities deliberately, where *that act itself* is accounted for by appealing to our credences and desires. In such cases, we choose to reason deliberately or to calculate expected utilities because doing so leads to greater expected utility than not doing so.¹⁴

Why do such cases arise? If we were ideally rational agents, such cases might not arise. If being ideally rational means always acting in ways that maximise expected utility without a need for deliberation or expected utility calculation, then engaging in such reasoning or calculation would just be a waste of time and effort, and might even decrease expected utility.¹⁵ But we are

¹⁴See Maher (1993), pp. 5-8, for a convincing defence of this point.

¹⁵There might also be an ideal agent who enjoys calculating expected utilities, or who is in a situation in which she has to calculate expected utilities on pain of being tortured. I set such cases aside.

far from ideal. When faced with hard decision problems, we might make choices that we will rue. If reasoning our way through a decision problem or calculating expected utilities yields greater expected utility than not doing so, we might well do so. Also, ideally rational agents might be ideally rational precisely because they are adept at deliberation and at expected utility calculation. In such a case, even such agents will engage in deliberate reasoning.

Calculating expected utilities or reasoning deliberately with various weights attached to the premises of an argument can of course be mentally taxing and time-consuming. Hence, engaging in such calculation or reasoning might decrease expected utility. So we sometimes try to simplify our calculations and reasoning. In some cases our credence in p may be high enough for us to assume that it is true, or, as Wedgwood puts it, to take it for granted and use it as a starting point in our argument. In some cases, when the probability of p occurring conditional on the performance of some act is negligible, we might not even consider the possibility of performing the act, and just factor it out of our expected utility calculations.¹⁶ But again, the fact that we do such things can be explained by appealing to our credences, in particular, to our credences regarding how likely it is that making such simplifications in our reasoning and calculations will lead to desirable outcomes. If making simplifications in our reasoning and calculations leads to greater expected utility than not doing so, then we might well make such simplifications.

If the above is correct, then the fact that we are sometimes ‘willing to use p as a premise in practical reasoning’, that we sometimes ‘simply take p for granted’, that we sometimes ‘treat p as a starting point for further reasoning’,

¹⁶Similar points have been made by Christopher S. Hill and Joshua Schechter (2007) and by Mark Sargent (2009) (pp. 116-118; pp. 245-247). Sargent, however, seems to think that because of the complexities of calculation and reasoning, decision theory itself sometimes dictates that we reason based on our binary beliefs, say, the belief that it is raining. But there is no need to appeal to binary beliefs—he can just hold that decision theory sometimes dictates that we reason based on certain unqualified premises, say, the premise that it is raining. See the following paragraph.

and that we sometimes calculate expected utilities with certain simplifying assumptions, can be explained without recourse to binary beliefs. That we are willing to use p as an unqualified premise in reasoning can be explained simply by saying that doing so yields greater expected utility than not doing so.

One might observe that there is some binary attitude involved when one is willing to use p as a premise in practical reasoning or when one takes p as a starting point in one's argument. Since neither credences nor desires are binary, one might think that this attitude is just that of binary belief, and that, after all, binary beliefs do some heavy lifting. But this reasoning is too quick: the attitude involved might well be *supposition* rather than binary belief. When we deliberate about what to do, our credence in $\neg p$ may be small enough so that to save time or to keep things simple, we just *suppose* that p is true. But once again, binary beliefs do not play any non-redundant role. The willingness to use p as a premise in reasoning is explained by the expected utility of *supposing* that p is true.¹⁷

One might also think that although practical reasoning can be accounted for by our credences and desires, there is room for binary beliefs to do some work. Keith Frankish claims that binary beliefs 'influence our behaviour in virtue of our partial beliefs and desires concerning them' (p. 88).¹⁸ He thinks that we reason with our binary beliefs and binary desires in conscious deliberation, although the decision to do what our binary beliefs and binary desires dictate is due to our credences and graded desires. Suppose we have a high enough credence that according to our binary beliefs and desires, we

¹⁷Note that nothing I have said in this paragraph should be taken as vindicating Lance. In simplifying calculations, one might suppose certain things that one is not committed to defending or that one does not accept. Also, Lance claims that decision theory is *conceptually dependent* upon the notion of acceptance, and what I have said here is compatible with the falsity of such a claim.

¹⁸He also thinks that beliefs must be deployable in deliberations that are truth-critical with respect to premises, i.e., deliberations in which we place high premium on using only premises that are true. I will say more about Frankish's view of binary beliefs in chapter 7.

should perform a certain act, and we have a great enough desire that we should do what our binary beliefs and desires dictate. Then performing the act will maximise expected utility.¹⁹

I have no quarrel with the view that we have credences and desires concerning what we believe and desire in a binary sense. And of course, a person who has high credence that she has a binary belief that p might act very differently from a person who has low credence that she has a binary belief that p . But similarly, a person who has high credence that there is a tiger in front of her will act differently from a person who has low credence that there is a tiger in front of her. Of course, tigers can influence our behaviour in virtue of our graded beliefs and desires concerning them, but this does not show that tigers do any heavy lifting in a way that is interesting to a proponent of folk psychology or decision theory. No doubt, in explaining and predicting action and behaviour, we have to refer to the contents of our credences and desires. But just because binary beliefs figure in such contents does not mean that they do any interesting work.

The point can be made stronger if we imagine a case in which a person has high credence that there's a tiger present and runs for dear life, when in fact, no tigers are present. The actual presence of a tiger is not essential to an explanation of the person's behaviour. Similarly, if, as is likely, people do not have perfect introspective abilities, we can imagine a case in which someone deliberates and then acts in a way that accords with her credences and graded

¹⁹Stalnaker (1984) seems to have a similar view regarding his notion of acceptance (pp. 91-92). According to him, 'Bayesian decision theory gives a complete account of how probability values, including high ones, ought to guide behaviour'. But then he notes that unless acceptance affects our behaviour, it is hard to see how the notion can be important for inquiry. His view is that 'the way one classifies [...] propositions as accepted, rejected, or judgement suspended makes a difference to how the agent behaves, or ought to behave' (*ibid.*). But if such classification is to be compatible with his view about Bayesian decision theory, then one's way of classifying propositions as accepted or otherwise would have to be understood in terms of the credences and desires that one has regarding the acceptance of those propositions. And then we can run an argument similar to the following to show that there is no non-redundant role for the notion of acceptance to play.

desires regarding her binary beliefs and desires, even though she does not have the binary beliefs and desires in question.

4.3.2 Theoretical reasoning

One might try to carve out a role for binary beliefs in our *theoretical* reasoning. After all, we do not just reason to help us decide what to do; we also reason to help us paint an accurate picture of the world. Even if binary beliefs play no role in the former, having a binary belief that p might dispose us to use p as a premise in our theoretical reasoning or to take p as a starting point in our theoretical arguments.

One might think that if we do not have a binary belief that p , then we might be unwilling to use p as a premise in theoretical reasoning even if doing so yields greater expected utility than not doing so. Conversely, one might think that if we have a binary belief that p , then we might be willing to use p as a premise in theoretical reasoning even if doing so yields smaller expected utility than not doing so. For one might think that theoretical reasoning is concerned with truth and not with increasing one's expected utilities.

But although theoretical reasoning might be concerned with truth, it is still an act in which we engage, and hence falls within the purview of decision theory. Suppose we know that if we use a particular proposition as a premise in theoretical reasoning, we'll move closer to the truth, but will be made to suffer physical pain. If we value truth more than the avoidance of physical pain, we will go ahead and use the proposition as a premise in theoretical reasoning, because doing so would lead to greater expected utility than not doing so.²⁰

²⁰Even if an act maximises expected *epistemic utility*, one will not perform it unless one values epistemic utility. See Levi (1980), pp. 50-51 and pp. 180-182, and chapter 6 of Maher (1993) for more on epistemic utility.

4.4 Binary Beliefs and Assertion

Yet another attempt to carve out a non-redundant role for binary beliefs seeks to tie them to our practice of assertion. If binary beliefs play a role in our acts of assertion that credences do not play, they contribute to some of the heavy lifting after all.

According to Kaplan (1996), ‘it is a mistake to suppose that we need recourse to talk about beliefs in order adequately to describe the doxastic input into rational decision making’ (p. 107). But he thinks that ‘talk of belief has a distinctive role to play [...] in epistemology and the philosophy of science’ that is tied to the practice of making assertions in the context of inquiry (*ibid.*). Kaplan recognises that we assert various things for various reasons. But he thinks that sometimes we aim only to assert the truth, and in such cases, we take people to believe what they are asserting. Hence, we have Kaplan’s Assertion View:

You count as believing p just if, were your sole aim to assert the truth (as it pertains to p), and your only options were to assert that p , assert that $\neg p$ or make neither assertion, you would prefer to assert that p . (p. 109)

According to Kaplan, the truth ‘is just an error-free, comprehensive story of the world’ (*ibid.*). As he points out, however, it is not reasonable to expect anyone to aim for a completely error-free and completely comprehensive story of the world. The more comprehensive the story is, the higher the risk of error. And the lower the risk of error, the less comprehensive the story. Some people might desire to be comprehensive more than they desire to avoid error; others might desire to avoid error more than they desire to be comprehensive. Hence, in the bid for comprehensiveness, someone who has a low credence (say, less than 0.5) in p might nevertheless prefer to assert p than to assert $\neg p$ or to

remain silent. On Kaplan's account, such a person counts as having a binary belief that p .

Kaplan ties binary beliefs to a sole aim (probably counterfactual) to assert the truth. He also ties one's preferences for asserting certain propositions to one's desire for comprehensiveness and one's desire to avoid error when trying to paint a picture of the world. But there is tension in Kaplan's position that binary beliefs are not essential to decision making, and yet, are intimately connected with assertion. For assertion is something that we do—it is an act that we may decide whether to carry out or not. One's aim to tell the truth and one's preferences regarding whether to assert certain propositions can be accounted for by one's credences and desires. If our sole aim is to tell the truth, then we desire to tell the truth more than we desire to do anything else. Also, the more we desire to avoid error, the less we desire to assert something we are not confident is true, but the more we desire comprehensiveness, the more we desire to assert it. In trying to account for why we would prefer to assert p or to assert $\neg p$ or to remain silent (given that these are our only three options), we do not need to appeal to binary beliefs.

4.5 Binary Beliefs and Epistemology

When we say that an agent's credences and desires allow us to account for the way she acts, we do not mean that they have to account for all aspects of it. Suppose she does something that is morally wrong, impolite, and aesthetically displeasing. It is no discredit to decision theory and to a folk psychological picture of beliefs and desires that they are silent on why the performed act has such qualities. Similarly, if the act of asserting that p or using p as an unqualified premise in reasoning is epistemically blameless (and not just pragmatically blameless), this aspect of the act need not be something

that the decision theory or folk psychological is committed to explaining.

However, given the important role that credences play in the explanation and prediction of action and behaviour, it's natural to wonder about their role in epistemology as well. In particular, it's natural to wonder if rational *credence* may take the place of rational *binary belief*.

One may think that the answer is 'No'. First, it seems that rational binary beliefs have a role to play in explaining why certain acts of reasoning and assertion are epistemically blameless. For example, it seems that while one rationally believes that p , one is epistemically blameless in asserting that p and in using p as a premise in reasoning, whereas if one merely has a high rational credence in p , one might not be so (more on this in the next chapter). Second, let us not forget that knowledge that p is traditionally thought to entail a rational binary belief that p .²¹ Insofar as we want to preserve the notion of knowledge, we might also want to preserve the notion of a rational binary belief.²²

But suppose that belief can be cashed out in terms of credences (and perhaps desires). This would vindicate Jeffrey's claim that Ramsey has sucked the marrow out of the binary notion of belief, since at the intentional level, credences (perhaps in conjunction with desires) would be able to play whatever role binary beliefs supposedly play in epistemology. In chapter 7, I shall argue for a way to understand the notion of binary belief that leaves it with no non-redundant role to play in epistemology. But before that, in the next two chapters, I shall consider various other attempts to understand how (rational) belief relates to (rational) credence, and discuss the problems that confront

²¹But dissenters include Radford (1966) and Lewis (1996).

²²One might also note that knowledge that p is thought to entail a *true* belief that p , and intuitively, it does not make sense to talk about credences being true or false. It is unclear, however, whether knowledge that p requires a particular doxastic attitude towards p that is itself true, or whether it requires the truth of p , and a particular doxastic attitude towards p , regardless of whether it makes sense to say that the doxastic attitude is true.

such attempts.

Chapter 5

Belief and Credence: A Dilemma

5.1 The Epistemic Notion of Binary Belief

How is belief related to credence? Answering this question will help us get clearer on the role of the latter in epistemology. In fact, if beliefs can be cashed out in terms of credences (perhaps in conjunction with desires), then whatever role that beliefs supposedly play in epistemology can be played by credences (perhaps in conjunction with desires) instead.

The main task in this chapter will be to set up a dilemma that confronts attempts to relate belief to credence. In the next chapter, I argue against certain attempts to resolve the dilemma, and in chapter 7, I defend my own attempt to resolve it. Henceforth in this chapter, *belief* will always mean binary belief, unless otherwise indicated.

Before setting up the dilemma, let us get clearer on the role of belief in epistemology. By the epistemic notion of belief, I mean a notion of belief that plays certain epistemological roles relating to knowledge, reasoning and assertion. This notion may come apart from the *ordinary* notion of belief, which has to do with what the word ‘belief’ means in ordinary language. The two notions may not diverge substantially, but if they do, then for my purposes, considerations about the epistemological roles that belief plays shall trump ordinary language considerations insofar as the investigation of how belief is

related to credence is concerned.

My own such investigation is premised on the principle that knowledge entails rational belief, as well as on the following two principles encountered in the previous chapter:

Belief-Reason If one rationally believes that p , one is epistemically blameless in using p as an unqualified premise in (deliberate and reflective) reasoning.¹

Belief-Assert If one rationally believes that p , one is epistemically blameless in asserting that p .

Note that even if we have a rational belief that p , the very act of using p as a premise in reasoning might make it no longer rational for us to have the belief in question. Suppose Ida rationally believes that she will never assert anything on Mondays. On Monday, asserting that she will never assert anything on Mondays is epistemically blameworthy, but this is not counterexample to Belief-Assert since the very act of making such an assertion ought to rob her of her belief. Also, suppose that Ida rationally believes that she is not a brain-in-a-vat (BIV). Her deliberating with the premise that she is not a BIV makes BIV possibilities salient, which may in turn lead her to believe that she does not know whether she is one. Suppose you think that it is epistemically blameworthy to use as a premise in deliberate reasoning what one believes one does not know. Then you will hold that it is epistemically blameworthy for Ida to use the premise that she is not a BIV in reasoning. But since it no longer seems rational for her to believe outright that she is not one, this is not a counterexample to Belief-Reason.

Some philosophers, such as Williamson (2000) and Keith DeRose (2002),

¹Weatherson (2005) is sympathetic to a similar principle which goes as follows: ‘If S justifiably believes that p , then S is justified in using p as a premise in practical reasoning’ (p. 419).

think that unless we *know* that p , we violate a norm in asserting that p . On such a view, even if we rationally believe that p , the assertion that p is in some sense improper if we do not also know that p . If such impropriety amounts to epistemic blameworthiness, then Belief-Reason and Belief-Assert are false (on the assumption that we can have rational belief without knowledge). One way to defend the two principles is to deny that knowledge is a norm of assertion.² But even those who think that knowledge is a norm of assertion will grant that a violation of a norm may be excusable or reasonable. For example, DeRose (2002) holds that someone who asserts that p without knowing it violates a norm, but is blameless if she reasonably thinks that she knows that p (p. 180). Likewise, Williamson (2000) holds that it is reasonable for one to assert that p if ‘one reasonably but falsely believes p , and is in no position to know that one does not know p ’ (p. 256).

However, DeRose and Williamson stop short of endorsing Belief-Assert. Williamson (2000) considers an example in which it is rational for us to believe that we shall not be knocked down by a bus tomorrow, but we know that we do not know (and presumably, are not reasonable in thinking that we know) that we shall not be knocked down (p. 255). According to Williamson, it would be ‘foolish’ to assert ‘baldly’ that we shall be not be knocked down by a bus tomorrow’ (*ibid.*). If he is right, then Belief-Assert is false (assuming that it is epistemically blameworthy to assert what one would be foolish to assert).

I agree with Williamson that in this example we ought not assert that we shall not be knocked down by a bus tomorrow. But I deny that the example is also one in which we rationally believe that p . If we ought not assert that p , either because we know that we do not know that p , or we do not reasonably think that we know that p , then it is simply not rational for us to believe that p outright. Williamson admits that such a denial is plausible when it comes to

²See, for example, Weiner (2005) and Lackey (2007).

occurrent beliefs; according to him, it is plausible that occurrent belief is some sort of inner assertion (pp. 255-256). But it is strange to hold that although it is not rational for us to believe occurrently that we shall not be knocked down by a bus tomorrow, it is rational for us to believe the same thing non-occurrently. No doubt, a non-occurrent belief that transforms into an occurrent belief may go from being rational to being irrational. Consider, for example, a non-occurrent belief that one won't entertain a thought involving monkeys today. When the belief transforms into an occurrent one, it becomes obvious that the belief is false. In such a case, even if it is rational for us to have the belief non-occurrently, it won't be rational for us to have it occurrently. But the belief about not being knocked down by a bus tomorrow is not like that—even when occurrent, it need not be obvious whether it is true or false.

None of the above provides *conclusive* support for Belief-Reason and Belief-Assert. But the principles are plausible enough that I shall base my investigation of the relationship between belief and credence upon their truth.

5.2 The Threshold View and the Lockean Thesis

According to Richard Foley (1993),

to say that you believe a proposition is just to say that you are sufficiently confident of its truth for your attitude to be one of belief. Then it is rational for you to believe a proposition just in case it is rational for you to have sufficiently high degree of confidence [credence, degree of belief] in it, sufficiently high to make your attitude toward it one of belief. (p. 140)

Foley terms the second claim *the Lockean Thesis*, since a hint of the claim can be traced to John Locke. Call the first claim *the Threshold View*, since it says that one's credence in p amounts to a belief that p by meeting a certain thresh-

old. Although I am primarily interested in the relationship between belief and credence, I shall also be discussing the relationship between rational belief and rational credence, since the two issues are interconnected. For example, if the Lockean Thesis is false, then so is the Threshold View, since the Threshold View entails the Lockean Thesis.

It is compatible with the Threshold View as stated that to believe that p is to have a credence of 1 in p . And it is compatible with the Lockean Thesis as stated that it is rational for one to believe that p just in case it is rational for one to have a credence of 1 in p . But as we shall see in the next section, such views lead to the unpalatable result that we have a rather restricted set of beliefs or rational beliefs. Yet if the threshold for belief or rational belief is less than 1, then proponents of either the Threshold View or the Lockean Thesis face a host of other worries, as I shall argue in section 5.4. In particular, they will have to contend with the problem of arbitrary thresholds, with the lottery paradox, with the failure of certain closure principles, and with the intuitive oddness of certain statements relating what one believes to what one is confident about. So a dilemma lurks. I shall look at ways to sharpen each horn of the dilemma in what follows, but note that the dilemma looms not only for those who subscribe to the Threshold View or to the Lockean Thesis. Even if, *pace* Foley, you think that talk of belief cannot be reduced to talk of credences, you might still think that to believe or to be rational to believe that p requires at least that one's credence in p be sufficiently high, while holding that whether a high credence is enough to give us belief or rational belief is another matter. In this case, you too will have to confront the dilemma.

First, a word about rationality. We are concerned with epistemic rationality, but epistemic rationality comes in different varieties. By 'rational', Foley means *egocentric rationality*, which roughly means *rationality by one's*

own lights.³ We can understand ‘rationality’ in other ways. For example, one might take it to refer to ideal rationality (whatever that is). Or one might understand ‘S is rational in believing that p ’ to imply that S is epistemically blameless in believing that p . One might also take it to be synonymous with ‘S is justified in believing that p ’—in fact some philosophers who discuss the relationship between belief and credence talk about justified beliefs rather than rational beliefs. Again we can understand ‘justified belief’ in different ways. Some take ‘S is justified in believing that p ’ to mean that S has good enough evidence to *know* that p (Fantl and McGrath, 2002). Others take it to mean that S is epistemically blameless in believing that p .⁴ Still others might take it to mean something different. Understanding ‘rationality’ or ‘justified belief’ in different ways will yield different versions of the Lockean Thesis. For my purposes, I shall read ‘It is rational for S to believe that p ’ as ‘It is epistemically blameless for S to believe that p ’, since it is the reading that goes most naturally with Belief-Reason and Belief-Assert.⁵ Similarly, I shall read ‘S is justified in believing that p ’ as ‘S is epistemically blameless in believing that p ’. Although different authors take ‘rational’ or ‘justified’ to mean different things, I shall suppress such differences when they do not matter and flag them only when they do.

5.3 The First Horn of the Dilemma: Belief and Full Credence

Suppose that belief or rational belief requires full credence. Then we will have a rather restricted set of beliefs or rational beliefs, since there are

³According to Foley, your egocentrically rational ‘belief must be able to stand up to your own critical reflection’, ‘must meet your own deep standards’ and ‘require that you have no grounds for dissatisfaction, given your current personal resources’ (p. 96).

⁴See Pryor (2001).

⁵Of course, one may hold that if we believe that p in an ideally rational way, or have good enough evidence for p to know that p , then we are epistemically blameless in asserting p or in using p as a premise in reasoning. But intuitively, this is because in such cases, it is epistemically blameless for us to believe that p .

not many propositions in which we have full credence or are rational to have full credence.⁶ First, if you hold that rational credence functions are regular, then no contingent proposition ought to be assigned a credence of 0 or 1. Or even if you hold the weaker view that we are not rational in having a credence of 1 in propositions whose truth we cannot establish merely from the armchair, then the only rational beliefs we have will concern either the *a priori* or our phenomenal experiences. Second, philosophers such as Kaplan (1996) and Maher (1986) have claimed that if we have full credence in p , then we should be prepared to bet life, limb and soul on the truth of p for a paltry return. But we will not bet our limbs—not to mention our lives or souls—on many propositions that we may ordinarily count as believing or as rationally believing. And such reluctance to bet seems perfectly in order.

The first horn of the dilemma is formidable. In the next chapter, I shall look at some attempts to blunt it, but let us now turn to the second horn.

5.4 The Second Horn of the Dilemma

Call a threshold that is less than 1 *non-extreme*. A non-extreme Threshold View or Lockean Thesis is one according to which the relevant threshold is non-extreme. Some of the arguments discussed below target the view that to believe or to be rational to believe that p , it is *sufficient* to have a credence in p that meets a certain non-extreme threshold. But as mentioned earlier, the dilemma confronts even those who reject this view, and hold on to the less controversial view that to believe or to be rational to believe that p , it is *necessary* to have a minimum credence in p that is non-extreme. Some of the arguments discussed below also purport to show that the second, weaker view faces certain difficulties. I do not think that these difficulties are insurmount-

⁶In the next chapter, we shall see how some philosophers, including Hawthorne and Weatherson (ms.), Pollock (2006), and van Fraassen (1995), attempt to tackle the first horn of the dilemma.

able (especially since I try to resolve the dilemma in chapter 7), but they are difficulties we have to confront nonetheless.

5.4.1 Arbitrary Thresholds

One worry is that any threshold that is sufficient for rational belief or is the minimum credence needed for rational belief is arbitrary, so long as it falls short of 1.⁷ Suppose the threshold for rationally believing p is 0.95. One might question why having a credence of 0.95 is sufficient for believing p , but having a credence of 0.949 (for example) is not. There are really two related worries here. First, it might seem incredible that such a significant difference between believing p and not believing p is really due to this minuscule difference in credence. Second, there might seem no reason to favour one value over another as the threshold.

One may respond by requiring that the threshold be vague, but this just shifts the worries. As Foley (1993) points out, vague thresholds may be no less arbitrary than precise thresholds (p. 142). Given that the threshold is vague, say, 0.95-ish, there remains the question why it should be 0.95-ish, as opposed to, say, 0.94-ish. Furthermore, as Weatherson (2005) notes, on either epistemic or supervaluational theories of vagueness, the worry about precise thresholds remains, since on such theories, there will still be a precise number such that one has a belief that p just in case one's credence in p is at least as great as that number (p. 421).

5.4.2 Descartes's injunction

The next worry targets the view that to believe or to be rational to believe that p , it is *sufficient* to have a credence in p that meets a certain

⁷Foley (1993), Stalnaker (1984), and Weatherson (2005) have all expressed similar worries (p. 142; p. 91; pp. 420-421).

non-extreme threshold. As Kaplan (1996) points out, the Threshold View makes nonsense of Descartes's injunction not to believe anything of which one is not certain (p. 98). For instance, if the Threshold View is correct, the relevant threshold is 0.9, and Jack's credence in p is 0.95, then he cannot help but believe that p . Descartes would then be asking for the impossible. (To add to Kaplan's point, it is also worth noting that if the threshold is 1, then the Threshold View makes Descartes's injunction pointless, since Jack simply would not be able to violate it.) But even if we disagree with Descartes's injunction, it does not seem incoherent (or pointless).

Kaplan does not discuss the Lockean Thesis, but at first glance, one might think that it fares slightly better than the Threshold View where Descartes's injunction is concerned. Even if it is false that one believes that p just in case one's credence in p is greater than or equal to 0.9 (say), it might still be true that it is rational for one to believe that p just in case it is rational for one to have a credence in p that is greater than or equal to 0.9. A proponent of the Lockean Thesis may just insist that Descartes's injunction is false, without holding that it is nonsensical. But suppose that the Lockean Thesis holds. Then to say that it is not rational for Jack to believe that p is just to say that it is not rational for Jack to have a credence in p that is greater than or equal to 0.9. Now it is reasonable to read the injunction as saying that it is not rational to believe that of which one is not certain. Then translating talk about rational belief into talk about rational credence, the injunction says that if Jack's credence in p falls short of 1, it is not rational for him to have a credence in p that is greater than or equal to 0.9. The Lockean Thesis renders Descartes's injunction absurd.

5.4.3 The Lottery Paradox

The third worry for the Threshold View and the Lockean Thesis is due to the lottery paradox, first discussed by Henry Kyburg (1970). In discussing the paradox, Kyburg appeals to a principle similar in essence to the Lockean Thesis to help him vitiate the closure principle that if one rationally believes that p_1 , that p_2 , \dots , and that p_n , then one rationally believes that $p_1 \wedge p_2 \dots \wedge p_n$. Of course, if one finds the closure principle independently compelling, one might take that as reason for why the Lockean Thesis is false.

But even if we set aside closure principles, the lottery paradox poses a problem for views like the Threshold View and the Lockean Thesis, according to which to believe or to be rational to believe that p , it is *sufficient* to have a credence in p that meets a certain non-extreme threshold. Consider an example in which there is an unbiased lottery with a million tickets, of which exactly one is the winning ticket.⁸ Suppose also that you have bought one of the tickets, and that you are apprised of all the relevant facts about the lottery. Your credence that any particular ticket is a loser should be very high (0.999999 in fact). But it is conceivable that you do not believe that the ticket you hold is a loser—it is conceivable that if you believe that a particular ticket is a loser, you will not have bought it in the first place. Note that the example works even if we change it so that your credence that your ticket is a loser is as high as we wish, so long as it is not 1. This is trouble for the (non-extreme) Threshold View, since having merely a very high credence in p is not sufficient for belief that p .

Does the above example pose a problem for the Lockean Thesis too?

Several philosophers who discuss the lottery paradox have the intuition that

⁸For the sake of simplicity, I have made some assumptions that are not needed for the example to work. First, the lottery need not be an unbiased one, and second there need not be a winning ticket. See Williamson (2000), p. 248; and DeRose (1996), p. 571.

you do not know that your ticket will lose. But among such philosophers, some have the intuition that you are rational in believing that your ticket will lose. Others have an intuition to the contrary—perhaps all you are rational in believing is the neighbouring claim that your ticket will *probably* lose. Dub this second intuition the lottery intuition.⁹

Mere head-butting of intuitions will just lead to a stalemate. Are there theoretical motivations for favouring one intuition over another? Someone who is sympathetic to the Lockean Thesis will of course be disposed to hold that it is rational for you to believe that your ticket is a loser. But the truth of the Lockean Thesis is precisely what is at issue here, so we cannot use it to support such a view. There is, however, an argument relating rational beliefs and practical reasoning to back up the lottery intuition. Hawthorne (2004), in arguing that one does not know that one's lottery ticket is a loser, asks us to consider the following:

The ticket is a loser.

So if I keep the ticket I will get nothing.

But if I sell the ticket I will get a penny.

So I'd better sell the ticket. (p. 29)

The above is an egregious piece of reasoning—it is epistemically blameworthy for you to employ the premise 'The ticket is a loser' in such an argument. This suggests that you do not know that your ticket is a loser, for if you had known, then the above piece of reasoning would have been appropriate. Although Hawthorne is concerned with knowledge, we can use his example to support the lottery intuition.

Hawthorne and Weatherson (ms.) write, '[I]t is a central part of

⁹For examples of the former, see Lewis (1996), p. 556; DeRose (1996), p. 573; and Foley (1993), p.164. For examples of the latter, see Ryan (1996), p. 126; Nelkin (2000), p. 376; Hawthorne and Weatherson (ms.), p. 8.

philosophical lore that if one believes that doing X will get you no money and doing Y will get you some money and you want money then, in the absence of competing desires, one will do X (unless one loses those beliefs and desires in the interim)?' (p. 8). If it is rational for you to believe that your ticket is a loser, then you are epistemically blameless in using the premise 'The ticket is a loser' in the argument above.¹⁰ And you will be epistemically blameless in concluding that you ought to sell the ticket. But suppose that the prize for the winning ticket is generous enough to undermine the conclusion (and that it is easy to see this). In such a case, you will not be epistemically blameless in reasoning as above, and so it will not be rational for you to believe that your ticket is a loser.¹¹ Notice that the argument works even if we change the example so that your credence that your ticket is a loser is as high as we like so long as it falls short of 1.¹²

In showing how the lottery paradox is a problem for the Threshold View and for the Lockean Thesis, I have not appealed to any closure principle. But I shall now examine some closure principles, and show how they contribute to the dilemma. As we shall see, closure principles threaten to fail if belief or rational belief does not require full credence. One might, however, think that the preface paradox (on which more later) provides independent reason to think that such principles fail anyway, and so their failure does not help

¹⁰Recall from the previous chapter that we do not need to appeal to rational beliefs to explain and predict action and behaviour—credences and desires do the job just fine. But sometimes we engage in reflective, deliberate reasoning to decide what to do—sometimes engaging in such reasoning yields greater expected utility than not doing so. And such reasoning admits of epistemic evaluation—if one is deemed rational in believing a premise, then one is deemed epistemically blameless in employing it in such reasoning.

¹¹Note that the kind of rationality at issue is epistemic rationality, even though the kind of reasoning involved is practical reasoning. Hawthorne wants to tie the epistemic notion of knowledge to practical reasoning. Similarly, we can try to tie the notion of epistemically rational beliefs to practical reasoning. For more on the relationship between knowledge and one's practical interests, see Fantl and McGrath (2002).

¹²*Pace* Hawthorne and Bovens (2000) who defend a version of the Lockean Thesis that has the following consequence: if the number of tickets in the lottery is big enough, there will be some tickets that the agent has to believe are losers on pain of irrationality (pp. 250-252).

sharpen the dilemma. But I shall address the worries that the preface paradox supposedly raises for such principles. I shall also provide direct arguments in support of the principles.

5.4.4 Closure

Consider the following closure principles:

Closure under pairwise conjunction For any agent S, and any p and q in S's set of beliefs, if S rationally believes that p and rationally believes that q , then S rationally believes that $p \wedge q$.

Closure under modus ponens For any agent S and any p and q in S's set of beliefs, if S rationally believes that p and rationally believes that q if p , then S rationally believes that q .¹³

Consider also the more general principle:

Closure under joint entailment For any agent S and any p_i and q in S's set of beliefs, if p_1, p_2, \dots , and p_n jointly entail q , and S rationally believes that p_1 , that p_2, \dots , and that p_n , then S rationally believes that q .

The third principle entails the first two principles. The first two principles are *prima facie* plausible, and the third principle seems plausible at least on the assumption that S is logically omniscient. Indeed, a commitment to closure under joint entailment or some principle like it might be partly responsible for the lottery intuition. If you rationally believe that Ticket 1 will lose, that Ticket 2 will lose, . . . , and that Ticket 1,000,000 will lose, and your beliefs jointly entail that there is no winning ticket, then by closure under joint entailment, it is rational for you to believe that there is no winning ticket. But it is irrational to believe such a thing. Assuming that you have no better or worse reason to

¹³I shall assume that modus ponens is a valid form of argument, *pace* McGee (1985).

believe that a particular ticket is a loser compared to any other ticket, it is not rational for you to believe of any ticket that it is a loser.

Is closure under joint entailment too strong? If p_1, p_2, \dots , and p_n jointly entail q , but we do not believe so, it might not be rational for us to believe that q . To get around this problem, we can modify the principle as follows:

Closure under believed joint entailment For any agent S and any p_i and q in S 's set of beliefs, if S rationally believes that p_1 , that p_2, \dots , and that p_n , and furthermore, rationally believes that p_1, p_2, \dots , and p_n jointly entail q , then S rationally believes that q .

But suppose for example that you rationally believe that p , that q and that p and q jointly entail $p \wedge q$, but you fail to put all your beliefs together. In such a case, it might not be rational for you to believe that $p \wedge q$. For instance, your mind might be fragmented, so that you rationally believe that p but not that q in one fragment of your mind, and that q but not that p in another fragment.¹⁴ Fair enough, but let's stipulate that the principle is meant to apply either to a single fragment of our mind, and to cases in which we have put all the relevant beliefs together. Or we may put forward the following principle, taking a leaf from Hawthorne's formulation of a similar principle involving the closure of knowledge:

Closure under competent deduction If S rationally believes that p_1 , that p_2, \dots , and that p_n , and furthermore, competently deduces that q while still rationally believing that p_1 , that p_2, \dots , and that p_n throughout, then S rationally believes that q .¹⁵

¹⁴For more on the fragmentation of the mind, see Stalnaker (1984), p. 83, and Lewis (1986), pp. 30-31.

¹⁵Hawthorne's (2004) principle is as follows: Necessarily, if S knows p_1, \dots, p_n , competently deduces q , while retaining knowledge of p_1, \dots, p_n throughout, then S knows q . (p. 33).

But for simplicity, I shall set aside the last two principles, and work with closure under joint entailment, closure under modus ponens, and closure under pairwise conjunction. Henceforth, when I mention ‘closure’ without mentioning which of the preceding three principles I’m referring to, I mean to refer to all three. I stipulate that closure is to be applied to an idealised agent—call her *Chloe*—whose credences satisfy the probability axioms, whose mind is not fragmented, and who performs deductions competently and infallibly.¹⁶ I shall also stipulate that closure is meant to apply to those cases in which, if Chloe rationally believes a proposition, she will not cease to do so after performing a deduction involving it. Note, however, that I’ve neither stipulated that Chloe’s credences always accord with the available evidence nor that she has perfect introspective access to her mental states nor that she is a perfect maximiser of expected utility—even the smartest logician may not always believe sensibly or know her heart’s desires or act according to them.

Closure is *prima facie* plausible insofar Chloe is concerned.¹⁷ Yet if rational belief does not require full credence, then it threatens to fail even for such agents.¹⁸ Suppose that p_1, p_2, \dots , and p_n jointly entail q , and that Chloe, who is the kind of idealised agent described above, has a credence of 0.95 in each of p_1, p_2, \dots , and p_n . And suppose that having a rational credence that is at least as great as 0.95 is sufficient for rational belief. All the probability axioms require of Chloe is that she assign a credence of no less than $1 - 0.05n$ to q , given that q is a contingent proposition.¹⁹ This means that so long as

¹⁶As Lasonen-Aarnio (2008) points out, an agent who is competent at deduction need not be infallible at it (p. 159).

¹⁷In fact, for Stalnaker (1984) and Lewis (1986), closure holds true for all our beliefs in a single fragment of our mind. If they are right, then it easily follows that closure holds true for all our *rational* beliefs in a single fragment of our mind.

¹⁸Many philosophers have argued that similar closure principles for knowledge fail when the (subjective or objective) probability of each of what one knows is less than 1. For example, see Williamson (2000), p. 117; Hawthorne (2004), p. 47; and Lasonen-Aarnio (2008), p. 159.

¹⁹The uncertainty of the conclusion of a valid argument is no greater than the sum of the uncertainties of its individual premises, where the uncertainty of a proposition is its

$n \geq 2$, the probability axioms allow her to assign a credence to q that is less than 0.95. For example, if $n = 11$, then Chloe's credence in q may well be less than 0.5, as far as satisfying the probability axioms is concerned. But then, it is not rational for her to believe that q . This example therefore poses a problem for someone who holds either the Threshold View or the Lockean Thesis, but who also subscribes to closure. But suppose that you reject the Threshold View and the Lockean Thesis, but maintain that one needs to assign a certain minimum credence to a proposition to be rational in believing it. It does not immediately follow that closure fails, for there may be some other conditions that rational beliefs have to satisfy, and that Chloe's beliefs may not have satisfied, that will help preserve the principle. If you subscribe to closure, the challenge for you is to lay down such conditions. (I attempt to meet the challenge in the next chapter.)

Note that closure is under threat even when very few beliefs are involved, so long as one does not need to have a credence of 1 in p to rationally believe that p . If Chloe rationally believes each of p and q with a credence of 0.95, and the threshold for rational belief is 0.95, the probability axioms alone allow that her credence in $p \wedge q$ be less than 0.95. We cannot conclude that Chloe rationally believes that $p \wedge q$, unless there is something in the notion of rational belief that together with the axioms ensures that Chloe's credence in q is no less than 0.95.

But why deal with idealised agents like Chloe instead of non-idealised agents like human beings? The relationship between rational beliefs and rational credences should be the same whether we are talking about human beings or about Chloe. For example, the lottery case shows that an agent's having a high (but non-extreme) rational credence in a particular ticket being a loser is not sufficient for her to be rational to believe that it is a loser—whether the probability subtracted from unity. See Adams (1998), p. 38.

agent is a human being or an idealised agent like Chloe is immaterial. Similarly, it will be strange if Chloe but not a typical human being needs to have full rational credence that it is raining (for example) to be rational to believe that it is raining.²⁰ But given that beliefs are related to credences in the same way whether one is talking about Chloe or about human beings, talking about Chloe has the advantage of making the discussion less messy. By idealising away from certain human limitations, we can help ourselves more readily to the tools provided by probability theory to discuss the relationship between credences and beliefs.

The Preface

One might think that since risk accumulates, it is not surprising that closure fails, and that even an idealised agent like Chloe is not rational in believing the conjunction of everything she believes individually. Some take this to be the lesson of the preface paradox, first discussed by David Makinson (1965). Supposedly, if Chloe is rational, she believes each single proposition she asserts in her non-fiction book. But supposedly, it would also be hubris for her to believe that the conjunction of what she asserts is true. In fact, it seems reasonable for her to add a preface to her book, granting that she has

²⁰Arguably, a human being who is told by a logic expert that a certain logical statement is true might be rational to believe, but not with absolute certainty, that the statement is true. But Chloe rationally believes the statement and also has full credence in it, since we have stipulated that her credences obey the probability axioms. Does this show that the relationship between Chloe's rational beliefs and rational credences is significantly different from the relationship between a typical human being's rational beliefs and rational credences, insofar as such beliefs and credences are about logical truths? No, for suppose that the threshold for rational belief is 0.95, and we stipulate that Ishban is an agent who has a rational credence of 0.99 in all true propositions that are of the kind: such-and-such is the capital of such-and-such a country. For example, if rationally believes that Canberra is the capital of Australia, then he has a rational credence of 0.99 in the proposition. But this holds because of what has been stipulated, and not because of any deep fact about the relationship between Ishban's beliefs and credences about the capitals of countries—if he had a rational credence of 0.97 in the proposition, he would count as rationally believing it.

probably made a mistake in (the non-prefatory part of) her book.²¹

In reply, one might insist that if Chloe is not rational in believing that the conjunction of the propositions asserted is true, then in the first place, she is not really rational in believing each individual proposition. According to Ryan (1991), it ‘seems extremely reasonable to assume that none of us who write books are ever actually justified [or rational] in believing [or asserting] each sentence in our books’ (p. 300).²² And perhaps, as Stalnaker (1984) suggests, there are some propositions that an author is only rational to believe with certain qualifications, say with the qualification that they are probably true. But why assert what she does not believe outright? Well, perhaps an author may be epistemically blameless even if she does not believe all that she asserts in her book. She just needs to *accept* that they are true, the way a scientist sometimes accepts certain idealising assumptions in order to further her inquiry. And this may all be perfectly reasonable—it may be reasonable to accept those things that we believe are false, if they are ‘*essentially* true, or close to the truth—as close as one can get, or as close as one needs to get for the purposes at hand’ (Stalnaker, 1984, p. 93; Stalnaker’s emphasis).

One might be unhappy with the above picture, and complain that there is still something epistemically blameworthy about Chloe if she asserts something that she does not believe outright. Let us grant that the complaint is right. But so what? This should not stop us, or Chloe, from writing books.

²¹One might hold that it is neither hubris for Chloe to believe the conjunction of what she asserts, nor to believe that she has made no mistake in writing her book. But chapter 3 of Christensen (2004) provides compelling reasons against holding such a position.

²²A potential worry with a position like Ryan’s is that it is just unintuitive that we are allowed so few rational beliefs in ordinary life. But how deeply such a worry runs depends on whether we are interested in the epistemic or the ordinary notion of belief. It might turn out that ‘belief’ as used in ordinary language just refers to high credence. In such a case, it will be worrying that we have very few rational beliefs, because it will be worrying that there are few propositions in which we are rational to have high credence. But it might also turn out that given the epistemic notion of belief, we have far fewer rational beliefs than we might have thought. Whether this is worrying depends on whether those rational beliefs that we have are enough to play the epistemic roles we want rational beliefs to play. I think so, but will only have the resources to address this issue adequately in chapter 7.

Chloe may be epistemically blameworthy in asserting things she does not believe outright, but in writing the book, she may be epistemically praiseworthy in presenting a comprehensive (and mostly correct) picture of the world. To invoke Kaplan (but with a slight twist), we can see the author as striving to maintain a balance between comprehensiveness and avoidance of error. In the bid for comprehensiveness, it will be hard to avoid error; so it will be hard in writing a comprehensive book to assert only what one believes. But the epistemic good in writing a comprehensive book may outweigh the epistemic sin of asserting what one does not believe outright.

Christensen (2004) offers an example in which we are supposed to have the intuition that a person is justified in believing each of a long list of claims, but is not justified in believing that the list is error-free (p. 35). He asks us to consider a person who looks up people's phone numbers in a directory, perhaps in preparation for a political campaign. He also asks us to imagine that the campaigner is very careful about checking that she got the phone number of each person right—she uses a ruler to line up each person's name with her number, and she checks each number two separate times before she records the data on her list. Christensen has the intuition that of each claim in the list, the campaigner is justified in believing that it is true, but since phone directories are not infallible, she is not justified in believing that all the claims are true. But such an intuition does not seem very stable to me. It seems that of any particular claim on the list, the campaigner may indeed be justified in believing that it is true *while* she is recording it on the list. But suppose the list has been completed. If the campaigner now picks a random claim from the list, while keeping in mind the fact that phone directories are fallible, there is at least some intuition that the campaigner is not justified in believing the claim outright. Rather than relying merely on intuitions about various preface cases, let us next examine some direct arguments for closure.

Deduction as a Rational Way to Increase Our Stock of Beliefs

I have addressed some worries about closure posed by the preface paradox. But what arguments are there in support of closure? Williamson (2000) points out that if deduction is a way by which we extend our knowledge, then we have reason to subscribe to certain principles involving the closure of knowledge (p. 117). One might think that, analogously, we can look upon deduction as a rational means by which we increase our stock of beliefs. The idea is that when Chloe reasons by modus ponens, for example, she is trying to get from her current belief that p and her current belief that q if p to a new belief that q . But if rational beliefs are not closed under modus ponens, then it might not be rational for her to increase her stock of beliefs in such a way, and it becomes harder to see the point of reasoning by modus ponens. We might also sometimes appeal to deductive arguments to help someone else increase her store of rational beliefs. But if closure fails, then we might not be able to do this. For example, Kaplan (1996) holds that if closure fails, then ‘the fact that we convince someone of the truth of each of the premises of a valid argument would seem to provide her no reason whatsoever to believe its conclusion’ (p. 97).

But the above argument will not convince someone who denies that deduction sets constraints on the way that rational beliefs get revised. Such a person may instead think that deduction sets constraints on the way rational *credences* get revised. For example, suppose Chloe’s credences in p and in $p \supset q$ are both 0.95. Reasoning by modus ponens will lead her to have a credence in q that is no less than 0.9.²³ To talk about such constraints on rational credences, there is no need to mention beliefs at all.

²³Proof: It is a theorem of probability that $P(A \vee B) = P(A) + P(B) - P(A \wedge B)$. Since $0 \leq P(A \vee B) \leq 1$, $P(A \wedge B) \geq P(A) + P(B) - 1$. Hence, substituting p for A and $p \supset q$ for B , $P(p \wedge (p \supset q)) \geq P(p) + P(p \supset q) - 1$, i.e., $P(p \wedge (p \supset q)) \geq 0.9$. But since $P(q) \geq P(p \wedge (p \supset q))$, $P(q) \geq 0.9$.

Furthermore, as Christensen (2004) points out, we may grant that deductive reasoning sets certain constraints on rational belief, but this need not mean that rational beliefs have to be closed (pp. 80-81). If rational belief is related to rational credence via the Lockean Thesis, for example, then deductive reasoning does set constraints on rational beliefs, but indirectly. If Chloe has a rational credence of 0.95 in p and in $p \supset q$, then her credence in q ought not be lower than 0.9. But if the Lockean Thesis holds, and the threshold for rational belief is 0.9, then Chloe ought to believe that q .

A related reason that Kaplan (1996) and Maher (1993) offer for closure points to the fact that we reason by *reductio ad absurdum*. If closure does not hold, why should Chloe conclude that p , if it is shown that the assumption that $\neg p$, together with various propositions that she rationally believes, leads to contradiction? For if closure under joint entailment fails, then even if her various beliefs entail that p , it need not be irrational for her to keep all her beliefs.

As Scott Sturgeon (2008) points out, however, ‘[t]he most glaring thing about reductio, in fact, is that its punch [...] is inversely proportional to the number of claims in use’ (p. 151). The more propositions we need to believe to infer that the assumption that $\neg p$ leads to contradiction, the less we are inclined to change our mind about p . Since reductio supposedly has a greater punch when the number of propositions needed to derive a contradiction is relatively small, the mere failure of closure under joint entailment does not render all instances of reasoning by reductio specious. Furthermore, we can think of a reductio as something that sets constraints on our credences, and not on our beliefs. As Ruth Weintraub (2001) puts it, a reductio ‘has the (rational) force to motivate us to change our mind when and only when it reveals to us that our credences are incoherent’ (p. 448).²⁴

²⁴See also Foley (1993).

Rational Beliefs, Deliberative Reasoning, and Assertion

Yet the notion of rational beliefs does seem to have something to with our epistemic practices, and one such practice is our engagement in assertion and in reflective, deliberative reasoning. I shall argue that this connection provides reason to think that Chloe's rational beliefs are closed under joint entailment. However, my argument will not appeal to how deductive reasoning is supposed to set constraints on rational beliefs. Rather, it will appeal to how rational beliefs help determine what makes for epistemically blameless reasoning and assertion in the first place.

Our stock of rational beliefs (in the same fragment of our mind) provides the propositions that we would be epistemically blameless in asserting or in deploying as premises in our reasoning. Now the stock might well get smaller throughout the reasoning process or throughout the series of assertions that we make. You may no longer believe your assertion that it is now exactly five o'clock after you have finished making the assertion. Or you may have an epiphany when you are reasoning, and decide to give up some of your beliefs. But consider what is left of the original stock of rational beliefs at the end of the reasoning process or series of assertions. You are epistemically blameless in having used p as a premise in your reasoning or in having asserted that p , so long as p is in the remaining stock of rational beliefs. Such a picture of rational beliefs provides support for closure. I will illustrate why by considering two simple cases.

Suppose that p , q , and r are in Chloe's stock of rational beliefs. Hence, she is epistemically blameless in using them as individual premises in her reasoning. But then she should also be epistemically blameless in using the premise $p \wedge q \wedge r$ in reasoning if doing so does not stop her from being epistemically blameless in using p , q , and r as individual premises.²⁵ For one thing, if

²⁵The qualification is meant to fend off certain counterexamples. For instance, if p is the

all three premises, p , q , and r , are true, then the premise $p \wedge q \wedge r$ will be true too. So if employing the first three premises in one's reasoning is not going to lead one awry, neither will the mere act of employing the latter premise.²⁶ But suppose closure fails. Then it is possible that $p \wedge q \wedge r$ is not in Chloe's stock of beliefs, in which case she may well be epistemically blameworthy in using $p \wedge q \wedge r$ as a premise in her reasoning, contrary to what has been argued.

Similarly, suppose that Chloe asserts that p , followed by q , followed by r . Since p , q , and r are in her stock of rational beliefs, she is epistemically blameless in asserting them. She should also be epistemically blameless in asserting $p \wedge q \wedge r$ if doing so does not stop her from being epistemically blameless in asserting p , q , and r individually.²⁷ In fact, it would be hard to make sense of why someone would be unwilling to assert that $p \wedge q \wedge r$, but would be perfectly happy to assert that p , q and r in the same breath (bracketing cases in which there are pragmatic considerations against asserting the conjunction, e.g. cases in which one is threatened with torture for making such an assertion, etc.).²⁸ Yet if closure fails, then $p \wedge q \wedge r$ might not be in Chloe's stock of rational beliefs, in which case she might not be epistemically

proposition 'I will not invoke conjunctions in my reasoning', one is epistemically blameless in using p , but not $p \wedge q \wedge r$, in one's reasoning. Using $p \wedge q \wedge r$ as a premise in one's reasoning stops one from being epistemically blameless in using p as an individual premise in that piece of reasoning.

²⁶Chloe is epistemically blameless in reasoning with the premise, v , if p , q , and r jointly entail v , but of course, she might not be epistemically blameless in reasoning with just the premise v , say, in a case in which p and q are enough to entail v , but to arrive at the conclusion u , one needs the premise r as well. Note that the entailment need not be obvious to us—Chloe, after all, is an idealised agent.

²⁷The qualification is meant to fend off certain counterexamples. For instance, if p is the proposition 'I will not assert any conjunctive statement', one is epistemically blameless asserting that p , but not that $p \wedge q \wedge r$. Asserting that $p \wedge q \wedge r$ stops one from being epistemically blameless in asserting that p .

²⁸Hawthorne and Weatherson (ms.) make a similar point (p. 15). Evnine (1999) hypothesises that it is part of our practice of belief attribution that if we attribute 'two propositions that are of no special logical or emotional significance to someone', then *ceteris paribus* (for example, when the beliefs are not in different fragments of her mind) we can also attribute to the person a belief in their conjunction (p. 214). Also, when a person explains her position about some subject matter, we can summarise her position by attributing to her one large conjunctive belief (*ibid.*).

blameless in asserting that $p \wedge q \wedge r$, contrary to what has been argued.

I shall go through three possible replies. First, one might hold that even if $p \wedge q \wedge r$ is not in Chloe's stock of rational beliefs, the mere fact that p , q and r are in the stock is enough for her to be epistemically blameless in employing $p \wedge q \wedge r$ in reasoning. But if she believes that p , that q , and that r , and if she is epistemically blameless in employing $p \wedge q \wedge r$ in reasoning, it is hard to see why we should deny that she rationally believes that $p \wedge q \wedge r$.

Second, one might reply that rationally believing each of p , q , and r without believing the conjunction only allows Chloe to be epistemically blameless in using each of them in a different piece of reasoning or in asserting them on different occasions—the assumption that Chloe is epistemically blameless in using any number of her rational beliefs as premises in the same piece of reasoning or that she is epistemically blameless in asserting any number of them in the same conversational context begs the question against the opponent of closure.²⁹ But we do not need to appeal to closure to support the assumption. Suppose Chloe has already employed p and q in a piece of reasoning and she is epistemically blameless in doing so. If she still rationally believes that r , then it is hard to see why she would be epistemically blameworthy in adding r as a premise to her reasoning. If she is indeed epistemically blameworthy in doing so, then she no longer has a rational belief that r anymore—she might have lost r from her stock of rational beliefs after deploying p and q in her reasoning. But this leaves intact the claim that at the end of your reasoning process, if a premise you have deployed remains in your store of rational beliefs, then you are epistemically blameless in having deployed it.

Third, one might hold that Chloe still rationally believes that r —it is just not deployable in reasoning. But it is dubious that we should count such an impotent belief as a proper belief. One might argue that Chloe's

²⁹Thanks to Leon Leontyev for pressing me on this point.

rational belief that r is not really impotent. Although she is not epistemically blameless in using r as a premise in her current piece of reasoning, she might be epistemically blameless in using it in a second piece of reasoning. But is this second piece of reasoning supposed to be one that takes place concurrently with the first piece of reasoning, or is it supposed to be one that takes place later? The former makes it hard to see why Chloe is epistemically blameless in using r in the second piece of reasoning, but not blameless in using it in the first piece of reasoning. The latter suggests that r is currently not in Chloe's stock of rational beliefs, but is allowed back in when she performs the second piece of reasoning. Either way, there is no reason to give up the claim that Chloe should be epistemically blameless in drawing from what is in her stock of rational beliefs to use as premises in her reasoning, so long as such rational beliefs are not jettisoned during her reasoning process. We can run a similar line of argument to show that Chloe is epistemically blameless in asserting any number of her rational beliefs in the same conversational context.

With the relationship between rational belief and epistemically blameless assertion in mind, let us return to the preface paradox. Christensen wants to hold that an author can be rational in believing each claim in her book even if she is not rational in believing, or is rational in disbelieving, the conjunction of the claims. But presumably, in writing her book, she is epistemically blameless in asserting all the individual claims in her book in one (long) breath, one after another. For if not, then there is reason to suspect that she is not rational in believing some of those claims after all. Now if she is epistemically blameless in asserting all the individual claims in her book in one long breath, then she is epistemically blameless in asserting the conjunction of the claims—from an epistemic point of view, we should not object to her replacing the full stops in her book with ampersands. But unless she rationally believes the conjunction of the claims, she might not be epistemically blameless in asserting it.

One might reply that the author is epistemically blameless in asserting each claim, although not all in the same breath. But since writing a book with a list of claims amounts to asserting all of them in the same breath, adopting such a reply requires one to say that the author is epistemically blameworthy in writing the book, thus diminishing the force of the preface paradox as an argument against closure. If an author is epistemically blameworthy even when she writes a book with claims each of which she rationally believes, it is hard to see what advantage there is in holding that she is rational to believe those claims rather than to hold, as Stalnaker does, that she merely accepts them. Other considerations such as those discussed in this section will then tilt the balance in favour of the proponent of closure.

Coarse-Grained Explanation of Action and Behaviour

How we use beliefs to provide a coarse-grained explanation of action and behaviour provides additional support for closure under joint entailment. To borrow an argument from Stalnaker (1984), suppose that Chloe believes that p (pp. 82-83). This is to say that she is in a state in which she is disposed to act in a way that typically serves her desires if p (together with her other beliefs) is true. Suppose also that Chloe believes that q . This is to say that she is in a state in which she is disposed to act in a way that typically serves her desires if q (together with her other beliefs) is true. It follows that Chloe is in a state in which she is disposed to act in a way that typically serves her desires if p and q , and hence if $p \wedge q$ (together with her other beliefs) are true. But this just means that she believes that $p \wedge q$.³⁰ Beliefs, and hence rational beliefs,

³⁰Note that the qualification in parentheses is important. As Stalnaker points out, when we appeal to a particular belief to explain our dispositions to act, such an appeal is accompanied by an appeal not just to our desires, but to other beliefs that we have as well. Weintraub (2001) argues that Stalnaker's line of reasoning fails, because '[a]n action can tend to be successful [...] in p -situations and in q -situations while tending to be unsuccessful in $p \wedge q$ -situations' (p. 445). To use her example, it is possible that a particular medicine will tend to cure someone if p is true and will tend to cure her if q is true, but will not cure her if $p \wedge q$ is

are closed under pairwise conjunction. It is easy to see how the argument can be extended to support closure under modus ponens and closure under joint entailment. (Remember the kind of idealised agent that Chloe is supposed to be.)

5.4.5 Belief and confidence

There is another way to sharpen the second horn of the dilemma without relying on whether closure under joint entailment is true for idealised agents like Chloe. Consider the following:

Iff

1. Jack believes that p if and only if q , and he's slightly more confident that p than that q .
2. Jack is 95% confident that p if and only if q , and he's slightly more confident that p than that q .

Identity

3. Jack believes that Mark Twain is Samuel Clemens, and he's slightly more confident that Mark Twain wrote *Huckleberry Finn* than that Samuel Clemens wrote it.
4. Jack is 95% confident that Mark Twain is Samuel Clemens, and he's slightly more confident that Mark Twain wrote *Huckleberry Finn* than

true. Cases like Weintraub's example are real and plentiful, but she has ignored Stalnaker's qualification about taking into account a subject's total set of beliefs. She quotes Stalnaker as saying that 'If a person is ... disposed to act in ways that would tend to be successful if p ... were true, and is also disposed to act in ways that would be successful if q ... were true, then he is disposed to act in ways that would be successful if $p \wedge q$ were true', but she omits the qualification 'together with his other beliefs' that comes after p , q , and $p \wedge q$. If a person is disposed to act in ways that would be successful if p together with her other beliefs (including her belief that q) is true, then she is disposed to act in ways that would be successful if $p \wedge q$ is true.

that Samuel Clemens wrote it.

Deterministic Causation

5. Jack believes that pressing the button will (deterministically) cause the world to end, and he's slightly less confident that the world will end than that the button will be pressed.
6. Jack is 95% confident that pressing the button will (deterministically) cause the world to end, and he's slightly less confident that the world will end than that the button will be pressed.

Conditional Reasoning

7. Jack is 20% confident that the world will end, but is 90% confident that it will end, conditional on someone pressing the button. On coming to believe that someone will press the button (and nothing more), his confidence that the world will end goes up to slightly less than 90%.
8. Jack is 20% confident that the world will end, but is 90% confident that it will end, conditional on someone pressing the button. On coming to be 95% confident that someone will press the button (and nothing more), his confidence that the world will end goes up to slightly less than 90%.

If you share my intuitions, there seems to be something odd about the odd-numbered sentences but not the even-numbered ones. Compare (3) and (4) for example. If Jack is slightly more confident that Mark Twain wrote *Huckleberry Finn* than he is that Samuel Clemens wrote it, then it does not seem rational for him to believe outright that Twain is Clemens. (Perhaps he cannot even have such a belief.) But he could well be rational in being 95% confident that Twain is Clemens. He might be 95% confident that Twain is Clemens on the basis of reliable testimony; yet, he might not be certain that they

are the same person precisely because he is just slightly more confident that Twain wrote *Huckleberry Finn* than that Clemens wrote it. Now let Jack's credence in 'Twain is Clemens' be as high as you like; so long as it is not 1, it seems that he is not rational in believing that Twain is Clemens. But this is problematic, since it does not seem that Jack needs to have full credence in 'Twain is Clemens' to be rational to believe it.

This intuition can be backed up by an appeal to how we invoke beliefs to provide coarse-grained explanations of action. For example, suppose that Jack accepts a challenge to place two different bets that will ensure that he breaks even. If Jack believes that Clemens is Twain, then presumably he will be disposed to place ten dollars on each of the following bets if they are available: one that pays him ten dollars if Clemens wrote *Huckleberry Finn*, and one that pays him ten dollars if Twain did not write *Huckleberry Finn*. But at the same time, if Jack is more confident that Twain wrote *Huckleberry Finn* than that Clemens wrote it, he is not disposed to bet in such a way.

Perhaps (3) does not sound too odd if we modify it to read:

3'. Jack believes that Mark Twain is Samuel Clemens, but is not certain of it; the reason is that he's slightly more confident that Mark Twain wrote *Huckleberry Finn* than that Samuel Clemens wrote it.

(3') does sound less odd than (3). But we may well be hedging when we say that Jack believes without being certain that Twain is Clemens. We may not actually want to say that Jack believes outright that Twain is Clemens, but rather that he believes that it is very likely that the two are identical. Also, there might indeed be a sense of 'belief' in ordinary language such that to believe that p is just to have a high credence in p . In fact, there are times when we use 'belief' to indicate uncertainty, as when we assert that we believe

so-and-so, instead of asserting so-and-so outright.³¹ If this sense of belief is operative in (3'), small wonder that (3') is felicitous.

However, we have seen that thinking of belief merely as high confidence leads to trouble: the view makes nonsense of Descartes's injunction, and leads to the failure of certain closure principles. Furthermore, there is still the task of explaining the oddity of (3) *as it stands*. This oddity suggests that there is another sense of 'belief' that cannot be cashed out solely in terms of sufficiently high credences. To recycle a point made in section 5.4.2, it is coherent to hold that because Jack is irrational to believe that Twain is Clemens if he is not certain of it. But it would be strange even to countenance the view that because Jack is not certain that Twain is Clemens, he should not be 95% confident that Twain is Clemens. So it seems that there is a sense of rational belief that cannot be cashed out solely in terms of sufficiently high rational credences. It is this sense of rational belief that is operative in the lottery case, generating the intuition that it is not rational to believe that one holds a losing ticket. It is this sense of rational belief that fits well with the closure principles discussed in the previous section. And it is this sense of rational belief that is epistemologically interesting, and that I attempt to explicate in the next chapter.

³¹As pointed out in Smithies (forthcoming).

Chapter 6

Resolving the Dilemma I

6.1 Seizing the First Horn of the Dilemma

The dilemma has reared its ugly head. How should we deal with it? In this chapter, I shall examine some responses to the dilemma, and argue that they are unsatisfactory.

Some philosophers have tried to seize the first horn of the dilemma. Levi (1980), as we have seen in chapter 4, holds that one needs to have full credence in p to believe that p (and hence, to rationally believe that p). I have argued in that chapter that Levi's view is problematic.

Hawthorne and Weatherson (ms.) have a view that bears some resemblance to Levi's. They hold that the threshold for belief is always 1. But they note that credences are always defined over a probability space, although which probability space they are defined over depends on context. They write:

When I am trying to determine [whether Professor Uncert will turn up for her talk tonight], should I consider the possibility that aliens have landed and destroyed the university, thus forcing the paper to be cancelled? [...] Obviously there is no requirement that I consciously consider [such a possibility]. What it seems is reasonable, and what it seems I actually do, is assign probabilities to some possibilities, while totally ignoring others. Given that I am finite,

and that there are at least beth-several possible worlds, it would be tricky for the possibility space over which my subjective probability function is defined to contain all the classes of worlds there are. So we should, and do, restrict the classes of possible worlds over which our subjective probabilities are distributed. [...]. The classes of worlds to which I assign positive subjective probabilities are [...] determined by contextual factors. (Hawthorne and Weatherson, ms., pp. 3-4)

The thesis that belief requires full credence seems implausible, because our having a credence of 1 in p suggests that we would be willing to bet life, limb, and soul on the truth of p .¹ But we are not prepared to bet like this on many of our beliefs. Hawthorne and Weatherson reply that such an objection ‘ignore[s] the contextual element to subjective probability’ (p. 5). Yes, we have many beliefs on which we are not prepared to bet life, limb, and soul. But according to them, this is compatible with the thesis that belief requires full credence. Suppose that one has full credence in p . The offer of a bet might shift the context in a way that makes certain $\neg p$ possibilities, previously non-salient, salient. Credence is then shifted to these possibilities, and since one no longer has full credence in p , one no longer believes that p .

But the thesis defended by Hawthorne and Weatherson is problematic. What does it mean to ignore a possibility on their view? According to them, one ignores p if one has ‘given no thought at all to whether or not p ’, and has a credence of 1 in a disjunction of possibilities inconsistent with p (pp. 3-4). But what does it mean to ‘have given no thought at all’ to p ? Most philosophers have given thought to whether we are brains-in-vats (BIVs). Does this mean that even when we are not currently thinking about BIVs, we cannot ignore

¹Here I assume that the very act of betting is independent of the truth of p (both objectively speaking and from our subjective perspective).

BIV possibilities? But if so, then even when we are not currently considering BIV possibilities, we do not believe (on the account offered by Hawthorne and Weatherson) that we have hands. Surely, this is not a consequence they intend their thesis to have.

So perhaps what they mean: ‘is *currently* giving no thought at all to p ’. Now if one has given thought to p , then even though one might not be giving thought to p currently, it seems that one may have a non-occurrent opinion about p . One may believe p , or disbelieve p . One may have a credence in p (that may well be zero). But on the view that Hawthorne and Weatherson advocate, so long as one has a credence of 1 in a disjunction of possibilities inconsistent with p , then if one is not giving any thought to p at the moment, we have to count her as ignoring p . Her credences would not be defined over p -possibilities. She does not even have a credence of 0 in p , since she has no opinion about p whatsoever.

This is implausible for reasons similar to those I offered against Levi’s view of belief in chapter 4. If Jack has thought about BIV possibilities before, it is plausible that he has an opinion about whether he is a BIV, even before the BIV possibility is raised again. Prior to the raising of such possibilities, he is typically disposed to act in certain ways. This disposition is best accounted for by holding that he has credence defined over BIV possibilities. For example, if he had been asked to bet that he was not a BIV, with nothing to gain if he wins, and having to suffer eternal torture if he loses, he might have been disposed to decline to bet. If he had been asked how confident he was that he was not a BIV, he might have been disposed to answer, ‘Very, but not absolutely, confident’.

Hawthorne and Weatherson will reply to the objection by holding that if Jack had been offered the above bet that he was not a BIV or if he had been asked how confident he was that he was not a BIV, the context would change.

BIV possibilities would become salient, and Jack's credences would become defined over them (pp. 6-7). We can then appeal to these newly-defined credences to account for Jack's dispositions to decline the bet and to answer that he is very but not absolutely confident that he is not a BIV. But such a rejoinder does not get to the heart of the matter. Jack's dispositions need not have been acquired upon being asked to bet or upon being asked the question of how confident he is that he is not a BIV—he could have had them all along. Admittedly, when the bet is offered or when the question is asked, the context changes. But what makes it true, prior to the raising of the BIV possibility, and before any change in context takes place, that Jack was disposed to reject any such bet and disposed to answer any such question in such a manner? It is hard to see what else would account for Jack's dispositions, except that he has some positive credence defined over BIV possibilities even before they become salient.

Bas van Fraassen (1995) also has a view of belief that is somewhat similar to Levi's, and Horacio Arló-Costa (2001) has recently defended a slightly modified version of the former's view. My complaint about van Fraassen's view is similar to my complaint about the views offered by Levi (1980) and by Hawthorne and Weatherson (ms.)—it has to do with the requirement that belief be accompanied by full credence. As we can see from the quote below, van Fraassen is comfortable with the requirement, and does not think that it leads to a paucity of beliefs. He writes:

I reject the naive Pascalian equation that a bet on p , with any payoff whatsoever, is worth to me my probability for p times that payoff. I think that Pascal's equation holds under restricted circumstances, with relevant assumptions kept fixed and in place. [. . .]. In a given context I have a number of full beliefs which delimit the presently contemplated range of possibilities; it is the latter which I grade

with respect to their comparative likelihood. The context may be anchored to a problem or type of problem, for which I go to this trouble. (1995, pp. 373-374)²

According to van Fraassen, we should reject the view that if our credence in p is 1, then even on a relatively small stake, we should be willing to bet life, limb, and soul on the truth of p .³ This is because according to him, a credence of 1 is spread over a delimited range of *presently contemplated* possibilities. But credences should not only be assigned to presently contemplated possibilities—as I mentioned in chapter 1, non-occurrent credences have an important role to play in the explanation and prediction of action and behaviour. Also, at the risk of repeating the same point *ad nauseam*, even while in a context in which only a delimited range of possibilities is contemplated, one might be disposed *not* to bet limb, life and soul on the truth of p . This disposition is best explained by saying that one does not have a credence of 1 in p (whether occurrently or non-occurrently) *even in that context*. Rejecting what van Fraassen calls the Pascalian equation leaves us unable to account for such a disposition.

A hard-line approach to dealing with the first horn of the dilemma insists that we have many beliefs or rational beliefs in which we have a credence of 1, and on which we are disposed to bet life, limb, and soul. According to Pollock (2006),

if you are warranted in believing p , then $P(p) = 1$. But then it follows that you can simply assume the truth of p in decision-theoretic reasoning. In particular, if you are offered a bet which

²I reject the equation too, because I reject the view that credences are just fair betting odds. But one can reject the view and hold that there is an intimate connection between one's credences and what one is disposed to do in various circumstances. In particular, one might think that if we have full credence in p and we're concerned only with winning a bet on the truth of p , then we are disposed to bet on p no matter the size of the stakes.

³According to van Fraassen, conditional probabilities are primitive, but what he says is also meant to apply to cases in which one has a credence of 1 in p conditional upon a tautology.

returns one cent if p is true and costs $\$x$ if p is false, that bet has a positive expected value regardless of what x is, and so you should accept it. (p. 108)

Although Pollock takes $P(p)$ to represent the objective probability that p rather than one's credence in p , someone sympathetic to what he says, and who also thinks that the probabilities involved in decision theory are credences, might hold that if one rationally believes that p , then one should accept the bet for any value of x . This seems like an incredible position to maintain, but Pollock tries to mitigate his position by claiming that '[w]e bet our lives on justified beliefs every day' (p. 109). For example, '[w]hen I take a bite of food I am betting that it has not been poisoned and that it is not infected with botulism', and when an airline pilot acts upon his belief that the plane has enough fuel, he 'is betting his life and the lives of his passengers' (pp. 109-110).

Pollock seems to think that the utility in living is infinite or unbounded. But so long as the utility of lives is finite, the relevant credences might be less than 1 but still be high enough so that the expected utility of taking a bite of food or of flying is higher than the expected utility of not taking a bite of food or not flying. I value my life, but do not assign it infinite utility. Also, just because I have some credence that something terribly bad will happen if I take a bite of food does not mean that the overall expected utility of taking a bite of food is lower than that of not doing so. For I may also have very high credence that I will die of starvation if I do not eat. In such a case, the utility of taking a bite of food and dying from poisoning may well cancel out the utility of not taking a bite of food and dying from starvation. Likewise, the utility of flying and dying in a plane crash may well cancel out the utility of not flying and dying in some other way.

One might worry that talking about utilities that are infinite or very large tends to distort our intuitions. Perhaps the mere threat of eternal torture

or death if p turns out to be false has a psychological effect of making us more cautious than necessary, even though we would have remained rational had we retained absolute certainty in p . But the first horn of the dilemma still glistens with menace even if we set aside talk about large, unrealistic stakes. Suppose that for some contingent p , we are absolutely certain that p , and we are rational in being so. Then we should be indifferent between betting \$20 on p and betting \$10 on p for a chance to win \$1. But I dare say that we ought not be indifferent. We ought to make the more conservative bet. For suppose it turns out that $\neg p$. Then it seems that we can be blamed for *risking* an extra \$10 unnecessarily. But if we had been rational in being absolutely certain that p , it would not have made sense to blame us for taking any such risk—from our point of view, there would have been no risk at all.

The first horn of the dilemma remains formidable. Let us now look at some attempts to deal with the second horn of the dilemma.

6.2 The Contextualised Threshold View

Given that the threshold for belief or rational belief is not in general 1, it is still possible that it is sometimes 1, and sometimes not. The *Contextualised Threshold View* says that to believe that p is to have a credence in p that meets a context-dependent threshold.⁴ Making the threshold context-dependent helps deal with the worry about having arbitrary thresholds for

⁴Does such a threshold vary with the context of the person making the belief attribution, the context of the subject to whom the belief is attributed, or the context of the person assessing the truth of the belief claim? Insofar as we want to maintain that a person's beliefs are intimately related to what she is epistemically blameless in asserting and to what premises she is epistemically blameless in employing in her reasoning, it would seem that whether she believes that p or not depends on the context she is in.

Note also that a similar issue arises with claims that 'know' is context-sensitive. One may think that the relevant context is the context of the attributor of the knowledge claim, the context of the subject to whom knowledge is attributed, or the context of the person assessing the truth of the claim. See MacFarlane (2005) for more on the differences between the three positions.

belief.⁵ By analogy, compare ‘belief’ with ‘tall’. Someone is tall if and only if her height meets a certain threshold. Now if this (possibly vague) threshold is fixed, then it might indeed seem arbitrary that it is fixed the way it is. But given that ‘tall’ is context-relative, this threshold is different in different contexts. Now one might still object that within each context, it is arbitrary that the threshold is fixed the way it is. But such an objection would be misplaced. Suppose, for example, that in the context of applying to become a model at a certain company, to be tall is to have a height of at least 1.85 metres. The requirement set by the company that its models be of such a minimum height may indeed involve some arbitrariness. But it’s not arbitrary to hold that within the context of such a requirement, one is tall if and only if one has a height of at least 1.85 metres.

The Contextualised Threshold View also seems to allow us to deal with the lottery paradox. Suppose that in a context in which one considers whether one will win the lottery, the threshold for belief is 1, but in normal non-lottery contexts, the threshold is less than 1. Then in the former context, if one’s credence in one’s losing the lottery is very high but imperfect, one does not believe that one will lose the lottery, even though one may believe in non-lottery contexts propositions in which one has imperfect credence.

Unfortunately, Kaplan’s point against the Threshold View works against the Contextualised Threshold View as well. Consider a context in which the threshold for believing that it will rain is 0.9. Someone whose credence meets that threshold in that context cannot help but believe that it will rain. If the Contextualised Threshold View is right, Descartes’s injunction that one ought not believe what one is not absolutely certain of would be rendered nonsensical in that context.

⁵I shall not repeat the point for other accounts of belief that take the minimum threshold for belief or rational belief to be context-dependent.

Also, the Contextualised Threshold View does not preserve closure for agents like Chloe, and does not explain the intuitive oddness of certain statements relating what one believes to what one is confident about. But I will postpone the discussion of these problems to when I discuss similar problems for the Contextualised Lockean Thesis.

6.3 The Contextualised Lockean Thesis

One may also put forward a contextualised version of the Lockean Thesis, and hold that it is rational to believe that p if and only if it is rational to have a credence in p that meets a context-dependent threshold. Suppose that in the lottery case, the threshold for you to be rational to believe that you will lose the lottery is 1. Then having merely a high rational credence that you will lose does not make it rational for you to believe that you will lose. But in normal contexts in which the threshold for rational belief is less than 1, you may well be rational in believing propositions in which you have high but not full credence.

But there is a problem. Consider the proposition that you have a lottery ticket. Your credence in this proposition might be lower than your credence that you will not win the lottery, but even in a lottery context it seems that you may rationally believe the former while it is not rational for you to believe the latter.^{6,7} This is suggested by the following example. Suppose that you are considering whether to throw away your lottery ticket. In such a case, you may be epistemically blameless in asserting that you hold a lottery ticket, but not in asserting that you will not win the lottery. (Assume that keeping the

⁶I owe this point to Alan Hájek.

⁷Let your credence function be represented by $Cr(\cdot)$. Suppose that $Cr(H)$, your credence that you have a lottery ticket, is 0.95, and suppose that $Cr(L|H)$, your credence that you will not win the lottery, given that you have a lottery ticket, is 0.999999. Since $Cr(L \& H) = Cr(L|H)Cr(H) = 0.94999905$, and $Cr(L \& \neg H) = Cr(L|\neg H)Cr(\neg H) = 0.05$, $Cr(L) = 0.99999905 > Cr(H)$.

ticket will cost you nothing.) But if the threshold for rational belief in such a context is 1, then it is not rational for you to believe that you are holding a lottery ticket after all.

There are other reasons why the Contextualised Lockean Thesis does not resolve the dilemma in a satisfactory manner. Consider the proposition that Twain is Clemens, in which it seems reasonable to have a credence of less than 1. But consider Jack, who is slightly more confident that Twain wrote *Huckleberry Finn* than that Clemens wrote it. It seems odd to say that Jack believes or rationally believes that Twain is Clemens, but what is the context which raises the threshold for believing or rationally believing such a proposition to 1? A natural answer is that it is one in which Jack has the occurrent thought that he is more confident that Twain wrote *Huckleberry Finn* than that Clemens wrote it. But even if Jack is not entertaining the thought consciously, it still seems odd to say, ‘Jack believes (or rationally believes) that Twain is Clemens, but is more confident that Twain wrote *Huckleberry Finn* than that Clemens wrote it’.

The Contextualised Lockean Thesis also does not help preserve closure for agents like Chloe. Suppose, for example, that Chloe’s credences in ‘If Sue presses the button, the alarm will sound’ and in ‘Sue presses the button’ are both 0.95, and it is rational for her to believe both claims. Insofar as satisfying the probability axioms is concerned, she might well be rational in having a credence of less than 0.95 in ‘The alarm will sound’. But if 0.95 is the minimum credence Chloe needs to have in a proposition to be rational to believe it, she does not rationally believe that the alarm will sound, *pace* closure. Holding that this is a case in which the context is such that the threshold for rational belief is 1 is implausible. Note that the example also poses a problem for the Contextualised Threshold View. Even if Chloe rationally believes that Sue will press the button, and that the alarm will sound if so, her credence in the

alarm sounding might well be less than the threshold for belief. In such a case, she does not believe, and hence does not rationally believe, that the alarm will sound.

None of the problems that plague the Contextualised Lockean Thesis are resolved if we hold that the threshold is vague.⁸ For in a lottery context, it does not seem rational to believe that one will lose the lottery unless one has a determinate credence of 1 that one will lose the lottery. But even in the same context, it may well be rational to believe that one holds a lottery ticket, even if one's credence that one holds a lottery ticket is vaguely 0.95 (say). Also, holding that the threshold is vague does not deal with why it sounds odd to say that Jack rationally believes that Twains is Clemens, but is determinately more confident that Twains wrote *Huckleberry Finn* than that Clemens wrote it. Finally, merely allowing the threshold to be vague does not preserve closure under joint entailment for agents like Chloe. Suppose the threshold for rational belief is vaguely 0.95. Suppose that Chloe's credence in 'If Sue presses the button, the alarm will sound' and in 'Sue presses the button' are both vaguely 0.95, and it is rational for her to believe both claims. It might well be rational for her to have a credence in 'The alarm will sound' that has a value clearly smaller than 0.95—say something in the interval between 0.91 and 0.93. In such a case, it is not rational for Chloe to believe that the alarm will sound.

6.4 Proposition-relative Thresholds

One might think that the threshold for rational belief is proposition-dependent rather than context-dependent. For example, on such a view, the threshold for rationally believing lottery propositions, such as the proposi-

⁸See Daniel Hunter (1996) and Scott Sturgeon (2008), who each defends a version of the Contextualised Lockean Thesis on which the threshold is vague. Both Hunter and Sturgeon are happy to give up closure.

tion that you will lose the lottery, is higher than the threshold for rationally believing other propositions, such as the proposition that you are wearing mismatched socks, or the proposition that you are holding a lottery ticket.

A related account of rational beliefs, advocated by Igor Douven (2002), holds that if one has a sufficiently high rational credence in p , and p does not belong to a set of propositions that is *probabilistically self-undermining*, then it is rational for one to believe that p (p. 396). According to Douven,

a set of propositions is probabilistically self-undermining relative to a person's belief state at a given time exactly if she at the time believes each proposition in the set to a degree exceeding t [the threshold] given her background knowledge at the time alone, but to a degree of t or less given her background knowledge plus [one] or more members of the same set. (*ibid.*)

Suppose we hold ticket₁, and there are n tickets all together, for some number n . Suppose also that there will be exactly one winning ticket. Our credence in our losing the lottery, given our background knowledge about the lottery set-up, might well exceed the threshold, t , needed to believe some non-lottery proposition, such as the proposition that we are wearing mismatched socks. But according to Douven's account, it is not rational for us to believe that we will lose the lottery, because our credence in our losing the lottery, given our background knowledge plus the propositions that ticket₂ is a loser, that ticket₃ is a loser, . . . , and that ticket _{n} is a loser, is 0, and hence less than t .

But neither account can do without an appeal to context if there is no fixed minimum credence needed for one to be rational to believe propositions such as the proposition that you are wearing mismatched socks or the proposition that you are holding a lottery ticket. Jonathan Vogel (1990) points out that many propositions, such as the proposition that the car that you parked

by your house has not been stolen, can be seen as lottery propositions.⁹ Out of the many people who park their cars by their houses, some will probably get their cars stolen. So parking your car by your house is like taking part in a lottery in which there is a chance that your ticket wins (your car gets stolen).¹⁰ It does not require too much ingenuity to concoct examples in which the proposition that one is holding a ticket and the proposition that one is wearing mismatched socks can also be seen as lottery propositions. Out of the many people who think that they have a lottery ticket, some will not actually have one; out of the many people who wear socks, some will wear mismatched socks. There is some chance that one is among those who do not actually have a ticket, and there is some chance that one is among those who wear mismatched socks. The threshold for rationally believing a proposition seems to be context-dependent rather than proposition-dependent after all. There will be contexts which are such that on Douven's account, even the proposition that one is wearing mismatched socks or the proposition that one owns a lottery ticket might be a member of a probabilistically self-undermining set of propositions.

Even if the two accounts considered above deal with the lottery paradox, they do not deal well generally with preserving closure for agents like Chloe or with claims relating one's beliefs to what one is confident about. For example, on such approaches, how do we maintain that if Chloe rationally believes that Sue will press the button, and rationally believes that if Sue presses the button, the alarm will sound, then she rationally believes that the alarm will sound? It is implausible that we need to have a credence of 1 in such propositions to be rational to believe them. And Douven's account does not prevent Chloe's credences in the first two propositions from exceeding t , but her credence in the

⁹See Vogel (1990), p. 22 and *passim*.

¹⁰But for a dissenting view, see Sharon Ryan (1996), pp. 134-136.

third proposition from falling below t . It is also doubtful that the two accounts can explain the oddity of the sentence ‘Jack believes that Twain is Clemens, but is more confident that Twain wrote Huckleberry Finn than that Clemens wrote it’. The proposition that Twain is Clemens hardly seems to be the kind of proposition in which one has to have a credence of 1 (assuming that ‘Twain’ and ‘Clemens’ are to be understood descriptively) to be rational to believe it. And on Douven’s account, if Jack’s credence in ‘Twain is Clemens’ exceeds t , and the proposition does not belong to a probabilistically self-undermining set, then Jack is rational to believe it. But the account is silent on why it sounds odd to add that Jack is more confident that Twain wrote Huckleberry Finn than that Clemens wrote it.

6.5 Beliefs, Acts, and Preferences

Instead of spelling out the relationship between credences and beliefs in terms of thresholds, one may try to develop a theory directly relating one’s beliefs to one’s preferences and to what one is disposed to do. One can then see whether the notion of belief that such a theory yields relates to credences correctly, and plays the epistemological roles we want it to play. Now, in giving a coarse-grained explanation of action and behaviour in terms of beliefs, the following principles are *prima facie* attractive:

Belief-Preference For all states of affairs and any agent S, S believes that p just in case S prefers as if p .

Belief-Act For all acts and any agent S, S believes that p just in case S is disposed to act as if p .¹¹

¹¹Versions of the principle have been proposed by several philosophers, including R. B. Braithwaite (1946), and (as we saw in the previous chapter) Stalnaker (1984).

What is it to prefer as if p ?¹² We can understand the locution in decision-theoretic terms. For all states of affairs r and s , one weakly prefers r to s if and only if one's expected utility in r is no less than one's expected utility in s , whereas one weakly prefers r to s given p if and only if one's expected utility in $r \wedge p$ is no less than one's expected utility in $s \wedge p$.¹³ S prefers as if p just in case for all states of affairs r and s , S (weakly) prefers r to s given p if and only if S (weakly) prefers r to s *tout court*. Suppose I prefer that you stay at home rather than go to the beach given that it will rain. If I do not believe that it will rain, I might prefer that you go to the beach rather than stay at home. But suppose that I believe that it will rain. Then, in accord with Belief-Preference, I will prefer that you stay at home rather than go to the beach: my unconditional preferences regarding whether you stay at home or go to the beach should be the same as my preferences regarding whether you stay at home or go to the beach conditional on it raining.

What is it to be disposed to act as if p ? For any act C and any other alternative act D , S is disposed to do C if she prefers a state of affairs in which she does C to one in which she does D , and S is disposed to do C given p if she prefers a state of affairs in which she does C and p is true to one in which she does D and p is true.¹⁴ S is disposed to act as if p just in case S is disposed to do whatever S is disposed to do given that p . Suppose that given my desires and beliefs, I prefer the state of affairs in which it rains and I carry an umbrella to one in which it rains and I do not. I am disposed to carry an umbrella given that it will rain. If I also believe that it will rain, I will prefer a state of affairs in which I carry an umbrella to one in which I do not *tout court*. In accord

¹²The provenance of the locution 'prefer as if' seems to be Fantl and McGrath (2002), although they focus on the locution 'rational to prefer as if'.

¹³I take one's preference for a state of affairs over another to be one's preference for a certain proposition rather than another to obtain. Note also that to say that S prefers r to s given p is *not* to say that if p , then S prefers r to s .

¹⁴Note that to say that one is disposed to do C given p is *not* to say that if p , then one is disposed to do C .

with Belief-Act, I will be disposed to carry an umbrella *tout court*.

At first glance, the principles are intuitively plausible. But they succumb to counterexamples. Suppose that I believe that it will rain and my credence that it will rain is 0.95. Suppose also that I know you are cash-strapped, and I want only the best for you. For a chance to win a hundred dollars, I prefer that you bet a million dollars on rain to not doing so given that it will rain. (Assume for such examples that utility is linear with money.) But since my credence that it will rain is 0.95, I do not prefer that you bet a million dollars on rain to not doing so *tout court*. So as it stands, Belief-Preference is false. It is also recognised that Belief-Act is vulnerable to counterexamples as it stands.¹⁵ Suppose I believe that it will rain and my credence that it will rain is 0.95. On a stake of ten dollars, I am disposed to bet a million dollars on rain given that it will rain. But since my credence that it will rain is 0.95, I will not be disposed to bet a million dollars on rain *tout court*.

Nonetheless, there seems to be something right about Belief-Preference and Belief-Act. Perhaps there are ways to modify the principles to preserve what is true about them, while avoiding the counterexamples discussed above. Some philosophers have offered restricted versions of the principles. I shall first go through their views, and then argue that they face certain problems.

6.5.1 Kyburg's version of Belief-Act

To avoid the kind of counterexample raised against Belief-Act and Belief-Preference earlier, one might try restricting the principles as follows:

Restricted Belief-Act For a *restricted* set of acts, options, stakes, states of affairs, etc., and for any agent S, S believes that p just in case S is disposed to act as if p .¹⁶

¹⁵See Kaplan (1996), pp. 104-105, and Isaac Levi and Sidney Morgenbesser (1964), pp. 222-223.

¹⁶Those who defend a version of this principle include Kyburg (1988), Nozick (1993), Fantl

Restricted Belief-Preference For a *restricted* set of acts, options, stakes, states of affairs, etc., and for any agent S, S believes that p just in case S prefers as if p .

We get different versions of the above principles depending on how we cash out ‘restricted’ and depending on what we restrict. For example, Henry Kyburg (1988) holds a version of Restricted Belief-Act:

[S] fully (r/s) believes [that p] just in case in any situation in which the ratio of risk of being wrong to the reward of being right is less than $r : s$, [S] simply acts as if [p] were true, i.e., refuses to bet against [p] at any odds. (p. 147)¹⁷

He gives an example to illustrate what he means:

Suppose I fully ($99/1$) believe that my car will start this afternoon, but I don’t fully r/s believe it for $r/s > 99$. Suppose the most that can be at stake in my acting as if the car will start is \$50.00. If a is an act that will cost me \$50 if my car doesn’t start, but otherwise will yield a benefit worth \$1, I will act as if I knew that the car would start. On the other hand, if we increase the ratio of the stakes to $200 : 1$, things change. If I now receive \$1.00 if the car will start, but lose \$200.00 if it doesn’t, I will perform the act in question exactly when the probability of the car starting is greater than $199/200$. (pp. 147-148)

In the first case, when the most that Kyburg stands to lose if the car does not start is \$50, any act equivalent to betting more than \$50 that the car will start

and McGrath (2002), and Dorit Ganson (2008), among others. I shall not discuss Ganson since she does not say how exactly the restriction is supposed to be imposed, other than that it is context-dependent (p. 452). Neither shall I discuss Nozick’s version of the principle, which is similar to Kyburg’s. But see Nozick, 1993, pp. 97-98.

¹⁷One might wonder if Kyburg thinks that there is just one notion of belief that he is analysing, or whether there are several of them, e.g. full ($99/1$) belief, full ($98/2$) belief, full ($97/3$) belief, etc.

(for a chance to win \$1) falls outside the restricted set of acts. Even if he will not bet \$200 that the car will start, he still counts as believing that the car will start, for out of the acts in the restricted set of acts, what he will do given that the car will start is just what he will do in fact. However, in the second case, even though his credence in the car starting remains the same as in the first case, he does not believe that the car will start. Given an expanded set of acts that includes the act of betting \$200 that the car will start, what he will do given that it will start is different from what he will do in fact.

6.5.2 Fantl and McGrath's version of Belief-Act

Fantl and McGrath (2002) defend a normative version of Restricted Belief-Act, although their aim is to advocate a view according to which what one counts as being justified to believe depends on certain pragmatic factors. To this end, they focus on defending a certain *necessary* condition for *justified belief*. According to them, 'S is justified in believing that p only if S is rational to act as if p ', where what is rational for S to do is understood as what maximises S's expected utility (p. 78). Fantl and McGrath (2002) think of expected utility as something that depends on the relevant epistemic probabilities, and not on the relevant credences (footnote 14, p. 79). But it is possible to read their principle in a way that links justified beliefs to credences. Does their principle commit them to holding that to be justified in believing that p , the probability (be it epistemic or subjective) that p is true has to be 1? No, for in a footnote, they say that they 'assume that an act (A) is rational for S just in case A is available to S and there is no available competitor B such that it is rational for S to prefer doing B to doing A ' (p. 73; see also footnote 11, p. 77). So more precisely, their version of Restricted Belief-Act can be stated as follows: for all *available* acts, if S is justified in believing that p , then S is pragmatically rational to act as if p . (I qualify the term 'rational' to avoid

confusion with the epistemic sense of ‘rational’ used to state the Threshold View and the Lockean Thesis.) Now suppose that I am justified in believing that it will rain and my credence or the epistemic probability that it will rain is 0.95. For a chance to win ten dollars, it is pragmatically rational for me to bet a million dollars that it will rain given that it will rain. But since my credence or the epistemic probability that it will rain is 0.95, it is not pragmatically rational for me to bet a million dollars on rain. This does not mean that I am not justified in believing that it will rain, since the act of betting a million dollars on rain might not be available to me. Fantl and McGrath do not say what exactly an available act is. But presumably, it is something of which one is aware. If there is an act that one may perform easily, and hence is in some sense available, but one is not even aware of the option of performing the act (and justifiably so), then the availability of the act is not going to affect what one is pragmatically rational to do, and will not determine what one is justified in believing.

6.5.3 Weatherson’s version of Belief-Preference

Weatherson (2005) offers a version of Restricted Belief-Preference. According to him, ‘an agent believes that p iff conditionalising on p doesn’t change any conditional preferences over things that matter’ (p. 422). More precisely, S believes that p just in case S (weakly) prefers A to B given q if and only if S (weakly) prefers A to B given $p \wedge q$, where A and B range over *live* and *salient* options, q ranges over *active* propositions, and S ’s credence in p is greater than $1/2$.¹⁸ A live option is an act that ‘is really possible for the agent to perform’, a salient option is one that the agent ‘takes seriously in deliberation’, and an active proposition q is a ‘(possibly degenerate) conjunction of propositions’ that

¹⁸He argues for the last clause in some detail, although I shall not delve into his argument. See Weatherson, 2005, p. 426.

is consistent with p , such that for each conjunct c , either the agent is ‘currently considering whether it is true’, or the agent is disposed to take seriously the question of whether it is true, and there are live and salient A and B for which ‘S weakly prefers A to B given c iff S weakly prefers A to B *tout court*’ is false (p. 423). Suppose, for example, that I believe that p . On Weatherson’s account, it is possible for me to believe that p even though my credence in p is less than 1, so long as for any active proposition q , my preferences given q are the same as my preferences given $p \wedge q$ (where my preferences range over states of affairs in which I perform acts that are live and salient).

6.5.4 Fantl and McGrath’s version of Belief-Preference

Fantl and McGrath offer a normative version of Restricted Belief-Preference. According to them ‘S is justified in believing that p only if, for any states of affairs r and s , if S is rational to prefer r to s , given p , then S is [pragmatically] rational to prefer r to s in fact’ (p. 77).¹⁹ Does their account require a credence or epistemic probability of 1 in p in order for S to be justified in believing that p ? No, for according to Fantl and McGrath, the principle is restricted to those cases in which the antecedent of the conditional is robust with respect to a situation in which S is faced with the choice of either making r true or making s true (pp. 74-75; footnote 8). In other words, the principle is restricted to those cases in which S would still be justified in believing that p , if S were faced with the choice of either making r true or making s true. Suppose that I am justified in believing that it will rain and my credence that it will rain is 0.95. On a stake of ten dollars, I am pragmatically rational to prefer betting a million dollars on rain given that it will rain to not betting a

¹⁹I’m concerned with seeing how the principle fares if treated as an explication of ‘justified in believing’, although Fantl and McGrath (2002) have a different concern. They are interested in showing that pragmatic considerations have a bearing on knowledge and on what one is justified in believing.

million dollars on rain given that it will rain. Nonetheless, since my credence that it will rain is 0.95, it is not pragmatically rational for me to prefer betting a million dollars on rain to not doing so *tout court*. But this is not a violation of the principle, since according to Fantl and McGrath, I would not remain justified in believing that it will rain, were I faced with the choice of either betting a million dollars on rain or not doing so.

6.5.5 Some Problems

How well do the restricted versions of Belief-Act and Belief-Preference considered above handle the dilemma? Suppose we own a lottery ticket but the act of throwing the lottery ticket away somehow falls outside the restricted set of acts—perhaps we do not take it seriously in deliberation, or it does not occur to us at all, or it is just somehow not available to us.²⁰ Then we might well believe or be rational to believe that our ticket is a loser even though our credence that it is the loser is less than 1: for all restricted acts and restricted states of affairs, we might be disposed to act as if the ticket is a losing ticket or prefer as if the ticket is loser. However, when our possession of the ticket is made salient to us or if the option of throwing the ticket away somehow becomes available, then we might no longer (be rational to) act as if or prefer as if the ticket is a loser. In such a case, we do not believe or it is not rational for us to believe that the ticket is a loser.

Some of the principles also vindicate certain versions of closure. Weatherston argues that his principle entails closure under pairwise conjunction for salient propositions, i.e., propositions that one consciously considers (pp. 426-427). Fantl and McGrath argue that if S is pragmatically rational to prefer as if p and pragmatically rational to prefer as if $p \supset q$, then S is pragmatically

²⁰It is hard to see, however, why the act of throwing away the lottery ticket is not available to us!

rational to prefer as if q (pp. 91-93). If we add to their argument the claim that for any p , S is justified in believing that p if S is pragmatically rational to prefer as if p , then what one is justified in believing is closed under modus ponens.

But more has to be said about available (or live) options. If an agent is not aware that an option is available, then the mere fact that it is available might not affect how she acts. (I would have taken the short-cut, but only had I been aware that it existed.) Also, an option may not be available, but the agent has high credence that it is. (If the option of a short-cut is not available, and I am mistakenly confident that it is, I may attempt to go for the short-cut.) So it seems that what is important to a defender of Restricted Belief-Preference or Restricted Belief-Act is not whether an option is really available, but whether the agent counts it as being available.

What is it for an agent to count an option as available? It is not enough for her merely to have some positive credence in it being so. Consider, for example, the proposition that it'll rain. Intuitively, we may believe the proposition even if our credence in the proposition is less than 1, and we have a small credence that the option of betting a million dollars that it'll rain for a possible return of one dollar is available. But if having *some* credence that an option is available is sufficient for it to be counted as such, then it will be much harder than one might think for us to have the belief in question. Even if our credence that it'll rain is very high, say, 0.999, we will not count as believing outright that it'll rain—we will prefer betting to not betting given that it'll rain, but prefer not betting to betting *tout court*.

Perhaps for an agent to count an option as available is for her to *believe* that it is. However, such a move appeals to the very notion that we want to cash out. It is also not plausible that an agent has to have a credence of 1 that an option is available for her to count it as such. If one is, say, 95% confident

that one has the option of betting a million dollars on rain, then it seems that the option should be considered to be available. Perhaps an agent counts an option as available if her credence in the option being available meets a certain high enough threshold. But then, one will have to say more about this threshold, and about whether some of the problems faced by the Threshold View and Lockean Thesis re-appear. For instance, there are questions about whether such a threshold is arbitrary and about whether it varies with context. There is also a question of whether the availability of options is closed under conjunction. If an agent takes option 1 to be available, option 2 to be available, ..., and option n to be available, does she also take *all* those options to be available? If n is large, the agent's credence in all those options being available may well be less than 0.5.

Even granting that the above issues are settled, there is still a problem with some of the principles we have encountered: they are overly restricted. There is a connection between our beliefs on the one hand, and what we are disposed to do in actual or hypothetical circumstances on the other hand. But restricting the set of options in the ways that have been surveyed leaves us unable to capture such a connection. Consider a case in which my credence that the car will start is 0.99, I believe and I'm rational in believing that the car will start, and the most that is at stake whether the car starts or not is \$50. Suppose that I know very well that the option of, say, selling a bet that will pay the buyer a million dollars if the car does not start and the seller a dollar if it does is *not* on the table, since I do not have a million dollars to begin with. Such an option is not available to me (using the terminology of Fantl and McGrath) and is not live (using Weatherson's terminology). According to the versions of Restricted Belief-Act propounded by Kyburg and by Fantl and McGrath, and according to the version of Restricted Belief-Preference propounded by Weatherson, the mere possibility of selling such a bet does

not affect whether I count as believing or as being justified in believing that the car will start. But suppose that I am asked what I would do in a purely hypothetical situation in which the option of selling the bet in question is live or available. The following piece of reasoning—label it *No-go*—would be plainly egregious:

The car will start.

If the car starts, I will earn one dollar by selling a bet that pays the seller a dollar if the car starts, and the buyer a million dollars if the car does not start.

So I will earn one dollar by selling a bet that pays the seller a dollar if the car starts, and the buyer a million dollars if the car does not start.

So if I want to earn a guaranteed dollar, I ought to sell such a bet if I have the means to do so.

While considering what I will do in the hypothetical situation in which the option of selling the bet is available or live, it seems that I neither believe nor rationally believe that the car will start. Otherwise, I might have reasoned, and given Belief-Reason, would have been epistemically blameless in reasoning as above.²¹ But since the option of selling the bet in question does not become available or live just because I am considering a hypothetical situation in which it is so, the theories under consideration all have to maintain—incorrectly—that I still believe and am rational in believing that the car will start.

No-go is a piece of practical reasoning. But examples involving theoretical reasoning also create trouble. Suppose that we are in an epistemology class, and someone questions whether it would be rational for me to believe that my car will start, pointing out that for all I am rational to believe, it

²¹Recall that according to Belief-Reason, while one rationally believes that p , one is epistemically blameless in using p as a premise in (deliberate and reflective) reasoning.

is possible that someone has sabotaged the car engine while I am attending the class. Suppose that my credence in my car starting remains the same, and it is still true that the most that I stand to lose if it does not start is \$50. Then according to the versions of Restricted Belief-Act propounded by Kyburg and by Fantl and McGrath, and according to the version of Restricted Belief-Preference propounded by Weatherson, the mere possibility of someone sabotaging the car engine while I'm in class does not affect whether I count as believing or as being justified in believing that the car will start. But suppose I let myself be drawn into the epistemological inquiry in which the class is participating, and I allow the bar for assertion and reasoning to be raised very high. Within the context of such epistemological inquiry, I am not epistemically blameless in asserting that my car will start or in using 'My car will start' as a premise in my reasoning. Then it is not rational for me to believe that the car will start. And if I am rational, I will cease to have such a belief.

Of course, using 'The car will start' as an unqualified premise in reasoning or asserting it may itself be counted as an option I can choose. So one might think that the above poses no trouble at all for the principles under consideration: for instance, in thinking about whether to sell the bet figuring in No-Go if one had the means to do so, the option of using 'The car will start' as a premise in No-go becomes available. Now I am not disposed and nor is it rational for me to be disposed to take up *this* option. Also, even if I prefer taking up the option to not doing so given that the car will start, I do not have such a preference *tout court*. One might claim that the theories actually predict that I do not believe or do not rationally believe that the car will start, while I'm considering whether to sell the bet in question.

However, Kyburg, as well as Fantl and McGrath, should be wary about counting the options of using certain premises in reasoning as available options. Remember that the most that is at stake in the above examples is \$50, whether

or not one is considering the bet in question and whether or not one is in epistemology class. (We can stipulate that merely considering the bet or being in epistemology class doesn't raise the stakes.) So for Kyburg, if the option of using 'The car will start' as a premise in No-go is relevant to whether one believes that the car will start when one is considering the bet in question (say), it should also be relevant when one is not doing so. Also, for Fantl and McGrath, if the option of using 'The car will start' as a premise in No-go is available when one is considering the bet in question, it is hard to see why it isn't available when one is not doing so. Intuitively, using 'The car will start' as a premise in No-go is something that one can easily do, unlike selling a bet with extremely high stakes. Now suppose that the option of using 'The car will start' as a premise in No-go is available or otherwise relevant to whether one believes or is rational to believe that the car will start, even when one is not considering the bet in question. In such a case, since one is not disposed and not rational in being disposed to act as if the car will start (by choosing the option), one does not count as believing or as being rational to believe that the car will start, contrary to how the above example above is set up. In fact, it will be hard to maintain that S believes or rationally believes that p for most mundane and contingent p . This is because the option of using p as an unqualified premise in some piece of reasoning like No-Go will usually be available, and typically, one would not be disposed or would not be rational in being disposed to choose the option.

According to Weatherson's version of Restricted Belief-Preference, the relevant set of restricted preferences range over salient options. Suppose that the option of using 'The car will start' as a premise in reasoning isn't salient when one isn't considering whether to sell the bet in question, or when one isn't considering the possibility that someone has sabotaged the car engine. Suppose also that the option becomes salient when one is considering such possibilities.

Then the above examples might not pose a problem for Weatherson. When the option isn't salient, one counts as believing that the car will start, because among the relevant options, one's preferences about which option to choose conditional upon the car starting are the same as one's preferences *tout court*. But when the option becomes salient, one no longer counts as believing the car will start. One prefers using 'The car will start' as a premise in reasoning to not doing so, conditional upon the car starting, but one does not have the preference *tout court*. Note, however, that in considering whether to sell the bet in question if I have the means to do so, I may not need to engage in deliberation before deciding that I would not sell the bet. In such a case, the option of using 'The car will start' as a premise in No-Go is not salient. It remains that I'm disposed not to use 'The car will start' as a premise in No-Go, and that it will be epistemically blameworthy for me to do so. But then Weatherson's version of Restricted Belief-Preference will not be able to explain why I believe that the car will start when not considering the bet, but lose the belief when considering it.

The criterion that an option be salient creates yet another difficulty for Weatherson. Consider a case in which a person is sleeping. Suppose that in a hypothetical situation in which she is awake (relative to the case that we are considering), she will refuse to place \$5 on a bet that returns \$10 dollars if p is true, but will gladly place \$6 on a bet of that returns \$10 if p is false. Presumably, even while she is asleep, it seems correct to hold that she does not believe that p . But if we cash out belief in terms of salient options, as Weatherson does, then since no option is salient to our subject while she is asleep, what are we to say about her beliefs? Perhaps in such a case, it is trivially true that she prefers to do what she would prefer to do on the assumption that p . Then on Weatherson's account, if her credence in $\neg p$ is greater than $1/2$, she counts as believing that $\neg p$. But there seems to be

something odd about ascribing a belief that $\neg p$ to someone when her credence in $\neg p$ is just slightly above 0.5. One would think that to have a belief in $\neg p$, one's credence in $\neg p$ should be close to 1.

Finally, what about the version of Restricted Belief-Preference offered by Fantl and McGrath? As it talks about preferences that range over states of affairs, and not just options that are live, available, or salient, it avoids some of the problems that I've raised against other versions of Restricted Belief-Act and Restricted Belief-Preference. But the principle is supposed to hold only if the following is true: were one faced with the choice of either making A obtain or making B obtain, one would remain justified in believing that p . If the principle is treated as an explication of what it is to be justified in believing that p , then it is not particularly illuminating: it purports to give us a necessary condition for one's being justified in believing that p , but the very stating of the necessary condition requires an appeal to the notion of *being justified in believing*.

I have considered various attempts to resolve the dilemma, and found them lacking. In the next section, I shall defend my own attempt to resolve it.

Chapter 7

Resolving the Dilemma II

7.1 Two Theses

We saw that if believing that p or rationally believing that p requires having full credence in p , then we will have very few beliefs or rational beliefs. But if full credence in p is not required, then we are faced with certain other problems—for instance, we will have to deny that rational beliefs are closed under joint entailment even for an agent like Chloe (our logically omniscient agent first introduced in chapter 5). I've argued that various attempts to resolve this dilemma are unsuccessful. But before defending my own attempt to resolve it, it is useful to note a parallel issue involving knowledge.

One way to think about knowledge is in terms of the elimination of possibilities: S knows that p if and only if S has eliminated all possibilities in which $\neg p$. But on the face of it there is a problem. There is little that we know if knowledge that p requires the elimination of *all* $\neg p$ possibilities. For example, it is arguable that we are unable to eliminate the possibility that while it appears to us that we have hands, we are actually handless brains-in-vats. If knowing that we have hands requires the elimination of such a possibility, then we don't know something as mundane as it. Perhaps knowledge that p does not require the elimination of all $\neg p$ possibilities. But one might follow Lewis (1996) in thinking that it 'just *sounds* contradictory' to hold that S

knows that p but cannot eliminate a possibility in which $\neg p$ (p. 549). To get around this problem, Lewis (1996) holds that although knowledge that p requires the elimination of all $\neg p$ possibilities, the quantifier ‘all’ ranges over a restricted set of possibilities, with context determining how this set looks. If he is right, we still know a fair amount in various contexts. For example, there are contexts in which ‘all’ is restricted to the set of possibilities in which we’re not brains-in-vats. We know that we have hands in such contexts if, within this restricted set, we have eliminated all possibilities in which we are handless.

I prescind from issues regarding whether Lewis’s account of knowledge is correct. We can nonetheless take a leaf from him, and hold that to believe or to rationally believe that p , one has to rule out all $\neg p$ possibilities within a restricted set of possibilities, with context determining how this set looks.¹ On this picture, it is consistent with S having a belief or a rational belief that p that some $\neg p$ possibilities outside the restricted set are not ruled out. (Just as on Lewis’s picture, it is consistent with S knowing that p that some $\neg p$ possibilities outside the restricted set of possibilities are not eliminated.) If to rule out a possibility is to assign it a credence of 0, we get the view that to believe or to rationally believe that p , one has to assign p a credence of 1 conditional on the proposition that carves out the restricted set of possibilities (where this proposition is picked out by context).² On this picture of belief or rational belief, we are not impaled on the first horn of the dilemma—when one believes or rationally believes that p , some credence may still be assigned to $\neg p$ possibilities outside the restricted set. Also, I shall argue in section 7.2 that we are not impaled on the second of the dilemma. (For instance, it will

¹I use the term ‘rule out’ rather than ‘eliminate’ when talking about belief, because on Lewis’s account of knowledge, ‘eliminate’ has a distinctive technical meaning. According to him, ‘a possibility W is uneliminated iff the subject’s perceptual experience and memory in W exactly match his perceptual experience and memory in actuality’ (p. 553).

²One may subscribe to a threshold view of ruling out on which to rule out a possibility is to assign it a low enough credence. But this will just lead one straight into the second horn of the dilemma.

turn out that within each context, Chloe's rational beliefs are closed.)

Let me now state the two main theses that I shall defend in this chapter more precisely. Let S be some arbitrary agent whose credence function and conditional credence function are represented by $Cr_S(\cdot)$ and $Cr_S(\cdot|\cdot)$ respectively. If S has a belief that p , the *context of inquiry* (of which more later) that S is in picks out what I shall call the *base proposition*—the proposition that carves out the set of possibilities within which S has to rule out all $\neg p$ possibilities to believe or to be rational to believe that p . Here are the theses:

Restricted Belief For any proposition p , any context of inquiry C , and any base proposition b , if S is in C and C picks out b for S , then ‘ S believes that p ’ is true if and only if and $Cr_S(p|b) = 1$.

Restricted Rational Belief For any proposition p , any context of inquiry C , and any base proposition b , if S is in C and C picks out b for S , then ‘It is rational for S to believe that p ’ is true if and only if it is rational for S to be such that C picks out b for S , and $Cr_S(p|b) = 1$.³

A few things to note. First, the above theses bear some resemblance to the accounts of belief offered by Levi (1980), van Fraassen (1995), and Hawthorne and Weatherson (ms.), each of which holds that belief requires full credence. I explained in the previous chapter why such a requirement is problematic; I maintain that a belief or a rational belief that p requires one to have a credence of 1 in p , *conditional* upon some base proposition b in which one has credence close to 1, where how close counts as close depends on the context of inquiry. Now since $Cr_S(b|b) = 1$, the requirement that $Cr_S(b)$ be close to 1 is important—without the clause, the theses will allow S to (rationally) believe a proposition in which she has low credence. Given the requirement, the theses

³The view of belief that I defend in this chapter is inspired by a blog post by Wolfgang Schwarz at <http://www.umsu.de/wo/2007/492>. I compare my view with some other similar views in section 7.3 of this chapter.

retain part of the spirit of the Threshold View and the Lockean Thesis.

Second, I shall say about more contexts of inquiry below, but as Restricted Belief and Restricted Rational Belief stand, they can be seen as analysis *schemata*—any debate about what a context of inquiry is or about how it picks out the base proposition can be seen as an internal debate among those who think that the theses are on the right track (schematically speaking). For example, just as those who agree to think about knowledge in terms of the elimination of possibilities might disagree about how context restricts the set of possibilities within which elimination takes place, those who agree that Restricted Belief and Restricted Rational Belief are on the right track (schematically speaking) might disagree about how the context of inquiry picks out the base proposition. And just as there is disagreement about whether it is the subject's or the knowledge attributor's context that determines whether one knows that p , there will be disagreement regarding the kind of context that determines whether one believes or is rational to believe that p .

Third, I won't spell out how exactly the base proposition is supposed to be determined by the context of inquiry. This puts me in no worse a position than Levi (1980) and Hawthorne and Weatherson (2004), who do not purport to say how exactly the agent's situation or context determines the set of serious possibilities (Levi), or the classes of worlds over which one's credences are distributed (Hawthorne and Weatherson). This also puts me in no worse a position than proponents of the Contextualised Threshold View or the Contextualised Lockean Thesis, who offer nothing more than perfunctory remarks on how the relevant thresholds are supposed to vary with context. The base proposition is just that proposition such that for one to believe or to rationally believe that p , one's credence in p conditional upon it has to be 1.

Fourth, as we shall see later, it will sometimes be indeterminate what the context of inquiry is. In such a case, it will be indeterminate what S be-

believes or is rational to believe. It will also be vague how high a credence one needs to have in a proposition for it to count as a base proposition. One might worry that because of such vagueness and indeterminacy, any attempt to cash out beliefs or rational beliefs in terms of base propositions is going to be hopeless. But as Lewis (1973) observes in defending his theory of counterfactuals, although the notion of comparative similarity is vague, it is precisely because of its vagueness that it is suitable for analysing counterfactuals, since it is sometimes vague whether a counterfactual is true (p. 91). Analogously, if it is sometimes indeterminate or vague what one believes or is rational to believe (as we shall see), it will be a good thing that it is sometimes indeterminate what the context of inquiry is, and sometimes vague how high a credence one needs to have in a proposition for it to count as a base proposition.

7.1.1 Whose Context?

I propose that a context of inquiry is one in which a subject is either deciding what to do, or is engaged in the search for truths. I take it that we often inquire into matters of both practical and theoretical import.⁴ In such contexts, we sometimes engage in reflective and deliberate reasoning (henceforth, just deliberation or reasoning) to help us reach a conclusion. But we typically *presuppose* certain propositions in deliberation, such as the proposition that we are not brains-in-vats or the proposition that we will not win the lottery, etc. Some of the presuppositions we make during deliberation will be relevant to our deliberation and others not. Suppose, for example, that in deliberating about whether to bring an umbrella to work, I presuppose that the umbrella will keep me dry if it rains and I also presuppose that tooth fairies don't exist. The first presupposition seems relevant, the second presupposition does not. Proposal: given that S is in a particular context of inquiry in which she has

⁴Cf. Kaplan's (1996) use of the term 'context of inquiry' (pp. 107-108).

certain beliefs or rational beliefs, there is a particular proposition in which S assigns a credence close to 1, and that is equivalent to the conjunction of relevant presuppositions that S makes or would make in deliberation. (Henceforth, unless otherwise stated, it should be taken as implicit that when we are talking about the propositions that S presupposes in a context of inquiry, we are concerned with those propositions that are relevant to S's deliberation.) In what follows, I shall say more about why it is the subject's context of inquiry that matters to what she believes or rationally believes. In section, 7.1.2, I shall say more about presupposition.

Suppose that Ida attributes to Jack a rational belief that he has enough flour at home to make a pizza, and Jack has a rational conditional credence of 1 that he has enough flour, given that his kitchen has not been ransacked by rats. Unbeknownst to Ida, Jack is wondering if it is possible that rats have ransacked his kitchen. He deems the possibility unlikely, but reasons that he should accept rather than reject a bet—call it the ‘no-loss bet’—that will cost him nothing if he has enough flour to make a pizza, but will reward him with a thousand dollars if not. For after all, he thinks to himself, there is a chance that rats have ransacked his kitchen. It seems that Ida is wrong—it is not rational for Jack to believe outright that he has enough flour. For suppose otherwise. Then given Belief-Reason (see chapter 5), in deliberating about whether he should accept the no-loss bet if it were offered to him, he would have been epistemically blameless in employing without qualification the premise that he has enough flour to make a pizza. But he would have been epistemically *blameworthy* in employing such a premise in his deliberation, and reasoning to the conclusion that he should be indifferent between accepting and rejecting the bet. The possibility that rats have ransacked his kitchen is relevant to whether he is rational to believe that he has enough flour, even though such a possibility is far from Ida's mind. The truth or falsity of Ida's belief attribution

depends on Jack's context of inquiry, and not Ida's.

Now suppose that in deliberating about whether to stop by the supermarket to buy flour, Jack is neither thinking about the no-loss bet nor considering the possibility that rats have ransacked his kitchen. But Ida has overheard some of Jack's friends talk about offering him such a bet. The mere fact that Ida is considering such a possibility does not mean that she would be correct in claiming that it is not rational for Jack to believe that he has enough flour. What is rational for Jack to believe does not depend on what Ida is thinking about, and Jack might well be epistemically blameless in using the premise that he has enough flour to make a pizza to reason to the conclusion that he need not stop by the supermarket. Once again, the truth or falsity of Ida's claim seems to depend on the features of Jack's context of inquiry.

Wedgwood (2008), who offers a contextualist account of *justified belief*, thinks that whether a person counts as being justified in believing that *p* sometimes depends on 'the standard that is favoured by a correct conception of one's [the attributor's] own situation' (p. 13).⁵ To use his example, suppose I ask you, 'Is anyone in the department justified in believing that Alice has accepted our job offer?' According to Wedgwood, if it is not that important to us whether or not Alice has accepted the offer, you may truly reply, 'John is justified in so believing', even if it is very important to John whether or not Alice has accepted the offer, and even if John speaks falsely when *he* utters, 'I'm justified in believing that Alice has accepted the job offer'.

But suppose that John is Alice's partner, and he intends to quit his job to fulfil his dream of becoming a writer, conditional on her accepting the offer. If she does not accept the offer, and he quits, they will not be able to pay

⁵He holds that sometimes it depends on the attributor's situation, whereas at other times, it depends on the subject's situation. Whether it is the subject's situation or the attributor's situation that matters in turn depends on whether the attributor's 'primary concern is with understanding [the subject or] with using [the subject] as a potential informant about whether *p* is true' (p. 13).

their bills, which would be disastrous. Suppose also that John believes that Alice has accepted the offer, and is thinking about whether to quit his job. According to Belief-Reason, if your claim that he is justified in believing that Alice has accepted the offer is true, John would be epistemically blameless in employing the premise ‘Alice has accepted the job offer’ in reasoning to the conclusion that he ought to quit his job. But he might not be epistemically blameless in doing such a thing if, given what is at stake, he ought not quit his job. Whether your claim is true does depend on John’s situation.

7.1.2 Presupposition

Intuitively, when one deliberates in a context of inquiry, one often makes certain presuppositions. For example, when Jack is wondering whether to drop by the supermarket to buy flour, he might presuppose that his kitchen has not been ransacked by rats. And such a presupposition seems relevant to whether he believes or rationally believes that he has enough flour to make a pizza. But what is it to *presuppose* a proposition? Since the term has been used to mean different things in philosophy, I shall tease out the notion that I am interested in. But I shall not provide an analysis of the notion—I know of none. I take it that the notion is an intuitive one, and shall let it take shape by contrasting it with other things that it is not.

First, I am interested in what *people* presuppose, and not in what *sentences* presuppose. It may well be that the sentence ‘The King of France is bald’ semantically presupposes that the King of France exists, so that if there is no such person, the sentence lacks a truth-value. But I am not interested in sentence presupposition.⁶

Second, I am interested in a notion of presupposition that is to be

⁶It might be that what sentences semantically presuppose depends on what people presuppose. For more on this, and on the distinction between the two kinds of presupposition, see Stalnaker (2002).

distinguished from *common ground*. If a proposition is part of the common ground, it is accepted by the participants in a conversational context, each of whom recognises that the other participants also accept it. Since one may disbelieve or have low credence in what one accepts, a proposition may be part of the common ground even if the participants in the conversation do not believe the proposition outright or have low credence that it is true.⁷ But I hold that whether a subject believes or rationally believes something depends on what the subject presupposes, and a subject may presuppose something even if it is not common ground, none of her conversational partners presupposes it, she believes that none of them presupposes it, all of them disbelieve it, or she believes that all of them disbelieve it, etc.

Third, one may presuppose a proposition even while entertaining its negation. Ida may be inclined not to factor counter-inductive and sceptical possibilities into her deliberation about everyday affairs even when such possibilities are salient, and she has positive credence in them. For example, it is physically possible that for the next ten minutes, all the air molecules in the room that Ida is in will gather at the ceiling, leaving her with no air to breathe. If you attempt to engage Ida in an epistemological inquiry by pointing out the possibility, and asking her what that shows about what she ought to believe, she may just brush off your attempt, and presuppose that such an unlikely scenario will not obtain. In so doing, she may believe and be rational in believing that she will not die of suffocation.

Fourth, propositions that are presupposed need not be considered or entertained consciously. In fact, entertaining such propositions consciously may sometimes result in your stopping to presuppose them or in your being no longer rational to presuppose them. Arguably, even a professed sceptic about

⁷See Stalnaker (2002) for more on the notion of common ground. Stalnaker takes *pre-supposition* to mean common ground.

the external world may go about her daily affairs presupposing that she is not a BIV. If she hurts her hand by accident, she may acquire the belief that her hand is in pain. But when BIV possibilities become salient to her, she may lose the belief because she stops presupposing that she is not a BIV. It might also be rational for her to believe that her hand is in pain before BIV possibilities become salient, but not after they become salient and she allows herself to be drawn into an epistemological inquiry with Cartesian standards in place.

A tricky case arises when the subject is sleeping or is otherwise not engaged in inquiry. It seems that it is rational for her to have various outright beliefs even while she is asleep, but how what determines the base proposition when she is not engaged in inquiry? I said earlier that we determine such a proposition by first determining the context of inquiry, and then by determining what S would presuppose when deliberating in such a context. In the examples considered so far, I have assumed that the subject concerned is actually engaged in inquiry. In such cases, her actual engagement determines what the context of inquiry is. But what if, say, the subject is sleeping?

I propose that in such cases the context of inquiry is determined by the inquiry in which the subject would be engaged were she engaged in one. For example, Jack might currently be asleep. But suppose that if he were awake and engaged in inquiry, he would be deciding whether to drop by the supermarket. It is rational for him to believe that there is flour in the cupboard: were he to deliberate about whether to drop by the supermarket, he would be rational in presupposing that rats have not ransacked the cupboard. Or suppose that if he were awake and engaged in inquiry, he would be thinking about whether to accept the no-loss bet. It is not rational for him to believe that there is flour in the cupboard: were he to deliberate about whether to accept the no-loss bet, he would not be rational in presupposing that rats have not ransacked the cupboard.

One might worry that the proposal does not sit easily with an ordinary conception of binary belief. The proposal allows that when Jack is asleep, it is indeterminate what he believes or rationally believes, since it is indeterminate what inquiry he would be engaged in were he awake and engaged in inquiry. But this accords with Belief-Reason and Belief-Assert. When Jack is asleep, would he be epistemically blameless were he awake and either asserting that there is flour in the cupboard or using the premise that there is flour in the cupboard in his reasoning? This question does not seem to have a determinate answer. Rather, it invites the further question: what is the context of inquiry in which Jack makes such an assertion or reasons in such a fashion? It is only if we have fixed the context of inquiry (and determined the base proposition that it picks out) that we can then proceed to find out what Jack believes or rationally believes in that context.

7.1.3 Some Implications

Given the theses and my account of how the base proposition is determined, practical considerations can bear upon rational beliefs. First, practical considerations can affect what one believes by affecting what one presupposes. Suppose I have a credence of 0.001 that my car engine has died, and this is the only reason I have a smidgen of doubt that the car will start. In a context of inquiry in which it will not prove too much of a bother to me if the car does not start, I might well presuppose that the engine is working. But in a context of inquiry in which I am deciding whether to sell a bet that pays the seller a dollar if the car starts, and the buyer a million dollars if the car does not start, I will not presuppose that the engine is working, at least if I am rational.

Second, practical considerations can affect what is rational for one to believe by affecting what is rational for one to presuppose. Suppose that in thinking about whether to sell the aforementioned bet, I presuppose that the

engine is working, and in so doing, count as believing that the car will start. Suppose that based on my belief, I then reason to the conclusion that I should sell the bet if I have the means to do so. But the conclusion is unreasonable—in such deliberation, it is not rational for me to presuppose that the engine is working, since the expected utility of selling the bet is lower than that of not doing so. Of course, the former will exceed the latter if my credence in the engine working is high enough. In such a case, it might well be rational for me to presuppose that the engine is working.

In some contexts of inquiry, however, it will be irrational to presuppose a proposition even if doing so does not lead one astray in practical deliberation. For example, suppose that I am engaged in a theoretical inquiry in which I am asking if it is rational for me to believe that the car will start. Suppose also that not much is at stake if it does not start. Whether it is rational for me to presuppose that the engine is working would depend on the epistemic standards set by the context of inquiry. If Cartesian standards are in place, then nothing in which one has less than full rational credence should be presupposed, and I should only believe those propositions in which I have full rational credence. If we loosen the standards ever so slightly, then to be rational to presuppose that the engine is working, my credence that it is will have to be very close to 1. If the standards are loosened just a bit more, then my credence that the engine is working will have to be fairly close to 1. And so on.

7.2 Applying the Theses

The theses allow us to avoid the first horn of the dilemma, by allowing us to believe that p or to be rational to believe that p without having full credence in p . But how do they fare with respect to the second horn of the dilemma?

7.2.1 Dealing with Descartes's Injunction

The theses do not make nonsense of Descartes's injunction to refrain from believing that of which one is not absolutely certain. Recall from chapter 5 that the Threshold View does not allow one to refrain from believing that p , if one's credence in p is above a certain threshold. Recall also that the Lockean Thesis renders Descartes' injunction absurd: if the threshold for rational belief is 0.95, for example, the injunction says that if one's credence in p falls short of 1, then one ought not have a credence in p that is greater than or equal to 0.95. But Restricted Belief allows people to refrain from believing that of which they are not certain. For example, consider a die-hard sceptic who refrains from presupposing anything in deliberation. She may enjoin other people to do likewise. Given Restricted Belief, this amounts to enjoining them to refrain from believing p if one's credence in p falls short of 1.⁸ Also, Restricted Rational Belief does not render Descartes's injunction absurd. Given the thesis, the injunction can be read as saying that if one's credence in p falls short of 1, one ought not presuppose that p .

7.2.2 Dealing with Lotteries

The theses also deal well with lottery propositions. Suppose that we have a lottery ticket, but in deliberating about whether to quit our jobs, we presuppose that we'll lose the lottery. Then by Restricted Belief, since our conditional credence in losing given that we'll lose is 1, we believe that we'll lose. And by Restricted Rational Belief, we are rational in doing so, if we are rational in presupposing that we'll lose, and we are rational in having a

⁸If absolutely nothing is presupposed, then for one to believe that p , one has to assign an unconditional credence of 0 to all $\neg p$ possibilities, where 'all' is unrestricted. Note that there is a sense in which you may choose what to believe, by choosing not to presuppose certain things. But there is also a sense in which you may not choose what to believe—given what you have presupposed, your credences determine what you believe, and you cannot will your credences to change.

credence close to 1 that we'll lose.⁹ However, in considering whether we'll win the lottery or whether to throw our lottery ticket away, we typically don't presuppose, and it is not rational for us to presuppose, that we'll lose. And since our credence that we'll lose conditional on whatever relevant proposition we have presupposed is less than 1, we don't believe outright, and it is not rational for us to believe outright, that we'll lose.

Note that in the same context of inquiry, we may presuppose a base proposition that entails that we are not brains-in-vats, or that appearance reflects reality, or even that we hold a lottery ticket. So if we have a credence of 1 in it appearing to us that we hold a lottery ticket, then our conditional credence in our holding a lottery ticket, given what we have presupposed, is 1. So we may believe or be rational to believe that we hold a lottery ticket, while not believing or not being rational to believe that we'll lose the lottery. This may be so even if our unconditional credence in the first is lower than our unconditional credence in the second.

7.2.3 Dealing with Closure

How do the theses deal with closure? Suppose that Chloe believes each of two or more premises that entail q , where q is contingent. Recall that given only the probability axioms, it is possible that even if Chloe's rational credence in each individual premise meets a certain threshold (less than 1), her rational credence in q falls short of it. The question is whether we can impose additional constraints on rational beliefs to preserve closure. I will argue that the theses provide such constraints.

Although the theses do not require Chloe to have an unconditional cre-

⁹Does this mean that cases may arise in which we believe p just because we've presupposed p and have sufficiently high credence in it? Yes, just as on Lewis's (1996) account of knowledge, one can know that one is not a BIV by properly ignoring BIV possibilities. It does not follow that it is easy to have rational beliefs, since it might not be easy to have high rational credence in what one presupposes.

dence of 1 in p to be rational to believe it, they do not lead to the failure of the closure principles. Or at least, the principles hold within a context of inquiry, even if they do not hold across contexts of inquiry. Let $Cr_C(.)$ represent Chloe's rational credence function and $Cr_C(.|.)$ her rational conditional credence function. Let b be the base proposition for Chloe in a particular context of inquiry. Suppose that p_1, p_2, \dots, p_n jointly entail q , and in that context, Chloe rationally believes that p_1 , that p_2, \dots , and that p_n . If $Cr_C(p_1|b) = 1$, $Cr_C(p_2|b) = 1, \dots$, and $Cr_C(p_n|b) = 1$, then $Cr_C(q|b) = 1$. Recall that by Restricted Rational Belief, Chloe's credence in b is high. If $Cr_C(b) = 0.95$ (say), then, since $Cr_C(q) \geq Cr_C(q \wedge b) = Cr_C(b)$, $Cr_C(q) \geq 0.95$. So closure under joint entailment (together with closure under modus ponens and closure under pairwise conjunction) is preserved—the thesis allow us to get around the worry that when there are many p_i s involved, each of which one rationally believes, one's credence in q threatens to fall below 0.5. A subject's rational credence in the base proposition sets a lower bound for the credence she needs to have in any other proposition to rationally believe it.

There might still be a threat against closure that comes from the *objective* risk that one's beliefs are false. To see why there might be a threat, let us first turn to Lasonen-Aarnio (2008), who argues that the following assumptions create trouble for the view that knowledge is closed (pp. 162-163).¹⁰ First, we know a lot of probabilistically independent propositions, each of which has some chance of being false. Second, to know that p , our belief that p should not be one that is susceptible to Gettier cases—it is not enough to have a belief that is true because of good epistemic luck. Third, if there is a low chance that p is true, then if it is true, it is lucky that it is. Now suppose that Chloe

¹⁰Lasonen-Aarnio (2008) states the closure principle that she is creating trouble for as follows: For all propositions p_1, \dots, p_n, q , and all subjects S, if S knows each of p_1, \dots, p_n , and S comes to believe q solely based on competent deduction from p_1, \dots, p_n , while retaining knowledge of each of p_1, \dots, p_n throughout, then S knows q (p. 158).

knows that p_1 , that p_2, \dots , and that p_n , and p_1, p_2, \dots , and p_n jointly entail q . If each of the p_i s has some chance of being false, the probability axioms alone allow a low chance that q is true when n is large. But if the chance that q is low, Chloe's belief that q would be lucky and would not amount to knowledge.

There is an analogous argument against the closure of rational beliefs. First, we rationally believe a lot of probabilistically independent propositions each of which has some chance of being false. Second, to rationally believe that p , it is not enough to have a belief that is true because of good epistemic luck. Third, if there is a high chance that p is false, then if p is true, it is lucky that it is true. Now suppose that Chloe rationally believes that p_1 , that p_2, \dots , and that p_n , and p_1, p_2, \dots , and p_n jointly entail q . If each of the p_i s has some chance of being false, the probability axioms alone allow a low chance that q is true when n is large. But then, Chloe's belief that q would be lucky and would not amount to a rational belief.

Now the mere fact that a true belief has a high chance of being false need not mean that it is not rational. Suppose that at a certain time, you hear from a usually reliable expert that a particular event will occur, and form a belief that it will occur. In fact the event occurs, but at that earlier time, the chance of the event happening was really slender—the expert uncharacteristically made some wrong calculations. So your true belief is lucky, but that might not stand in the way of its being rational. Of course, if one's standard for having a rational belief that p is to have evidence for p that is good enough for one to *know* that p , then the fact that the belief has a high chance of being false prevents it from being rational if it also prevents you from knowing that p . Also, suppose we replace talk about chance with talk about some kind of epistemic probability, and we hold that rational belief requires high epistemic probability. Suppose also that you believe that p , p happens to be true, and the relevant epistemic probability that p is true is very low. Your belief that p seems lucky, and does

not seem rational.

But given Restricted Rational Belief, there are two natural conditions to impose on rational beliefs to help preserve closure. First, hold that a sufficiently high rational credence in the base proposition, b , is one that is accompanied by a sufficiently high epistemic probability that b is true, high enough not to make a belief that b lucky. Second, hold that for any p_i that is rationally believed, one's rational credence in p_i given b is 1 only if the relevant epistemic probability that p_i given b is also 1. Given these two conditions, the relevant epistemic probability that q is true given that b is true will be 1, and the relevant epistemic probability that q is true will be no less than the relevant epistemic probability that b is true. So the relevant epistemic probability that q is true will be high enough not to make the belief that q lucky, even though the relevant epistemic probability that any particular p_i is true falls short of 1.

7.2.4 Dealing with Confidence Claims

Let Jack's conditional credence function be $Cr_J(.|.)$ and b be the base proposition. For Jack to be rational to believe that Twain is Clemens, $Cr_J(\text{Twain is Clemens}|b) = 1$. But then, the following ought to hold: $Cr_J(\text{Twain wrote } \textit{Huckleberry Finn}|b) = Cr_J(\text{Clemens wrote } \textit{Huckleberry Finn}|b)$. Now this does not quite explain the oddity of:

3. Jack believes that Mark Twain is Samuel Clemens, and he's slightly more confident that Twain wrote *Huckleberry Finn* than that Clemens wrote it.

If Jack believes that Twain is Clemens, then it seems that by the theses, all we may conclude is that conditional upon b , Jack has to be as confident that Twain wrote *Huckleberry Finn* as he is that Clemens wrote it.

But if in believing that p , we presuppose some special proposition, it is plausible that when we employ the words ‘confidence’ and ‘certainty’ in the same breath as we employ ‘belief’, we also presuppose the same proposition. It is plausible that we read the second conjunct of (3) as saying that Jack has higher credence that Twain wrote *Huckleberry Finn* than that Clemens wrote it *conditional on b* (where the ‘conditional on b ’ bit may well be tacit). But this should not be so, since Jack’s credence in Twain not being Clemens conditional on b is 0.

Admittedly, more has to be said to defend the above reading of (3). If the reading is correct, then Restricted Belief and Restricted Rational Belief at least have the resources to explain the oddity of (3), whereas other accounts of the relationship between (rational) belief and (rational) credence that stop short of holding that (rational) belief requires full credence are not able to do so.

7.3 A Comparison with Similar Views

My view about how (rational) belief and (rational) credence are related bears some resemblance to Levi’s view (discussed in chapter 4), and with the view put forward by Hawthorne and Weatherson (discussed in the previous chapter)—the crucial difference between their views and mine is that they hold that belief requires full credence, whereas I hold that it requires a conditional credence of 1.

Another view similar to mine has been defended by Jonathan Roorda (1995) in an unpublished manuscript.¹¹ Roorda distinguishes between *strong belief* and *weak belief*. To believe something strongly is to have an unconditional credence of 1 in it. To believe something weakly, one does not need to have an unconditional credence of 1 in it. But one needs to have a credence

¹¹The manuscript is available at: <http://www.princeton.edu/~bayesway/pt/Wolfman.pdf>.

of 1 in p conditional on what is presupposed in one's conversational context. Roorda, however, means by 'presupposition' what Stalnaker (1973) and Lewis (1979) take it to mean, namely, common ground. In this sense of the term, we may presuppose something that we find 'highly doubtful' or that we 'even disbelieve [...] altogether'. Unfortunately, this means that one can have a weak belief in a proposition in which one has low unconditional credence. But intuitively, weak beliefs are *not* beliefs. Furthermore, such a notion of belief does not play the role of being the doxastic component of knowledge, insofar as we think that to know that p , one needs to have a high enough credence in p .

There is also a problem with another requirement that Roorda imposes upon weak beliefs. He holds that for him to have a weak belief in p , 'there is no probabilistic judgement with which [he] identif[ies] conditional on [p] but with which [he does] not identify unconditionally' (p. 22). For example, his opinion about whether a flat federal income tax will be adopted conditional on a Clinton or Dole victory has to be the same as his opinion about whether such a tax will be adopted *tout court* for him to count as weakly believing that Clinton or Dole will win. But this will mean that he does not believe that Clinton or Dole will win, if his unconditional credence in Clinton or Dole winning the election is less than 1. Suppose he judges that it is probable to degree 0.9 that a flat federal income tax will be adopted *and* Clinton or Dole will win the election, conditional upon Clinton or Dole winning the election. But he would not presumably judge that it is probable to the same degree that the conjunction holds true *tout court*. Given that we do not have a credence of 1 in most kinds of proposition that we presuppose (in Roorda's sense of the word), the set of propositions that we weakly believe becomes rather restricted by Roorda's own lights.

My view is also very similar to one outlined by Wolfgang Schwarz

(2007). Schwarz's view differs from mine in that he holds that it is the belief attributor's context, rather than the subject's context, that determines whether the subject believes something or not. In section 7.1.1, I explained why it's the subject's context rather than the belief attributor's context that matters. Schwarz also allows that the subject has low credence in the proposition that carves out the set of possibilities within which she has to rule out all $\neg p$ possibilities in order to believe that p . In discussing Roorda's view, I explained why this is problematic.

Finally, John C. Harsanyi (1985) has a view about *acceptance* that bears some resemblance to my view about *belief*. In certain cases in which a subject is deliberating about what to do, it will make things easier for her if she simplifies the decision problem that confronts her by treating some proposition, p^* , as if she has a credence of 1 in it, i.e., by being practically certain (using Harsanyi's terminology) that p^* is true. Then instead of working with unconditional probabilities, she can work with probabilities conditional upon p^* being true. If the probability of a proposition conditional upon p^* is 1, then she also counts as being practically certain of that proposition. Now such decision problems may involve practical matters such as whether to drop by at the supermarket after work, or they may involve theoretical matters, such as those involving which scientific theory to adopt—the important thing to note is that simplifying our decision problems can lead to greater expected utility than not doing so (p. 5; p. 14). As an illustration, if Jack is deliberating about whether to drop by the supermarket to get flour, it might be easier for him if he works with his probabilities conditional on no rats having ransacked his kitchen, if his credence in rats having ransacked his kitchen is low enough.

One can simplify one's deliberation even more by being practically certain of p^* in a persistent way, i.e., by treating p^* as if one has a credence of 1 in it in all decision problems that one encounters (p. 8). The effort that might

be involved in checking whether one is justified in being practically certain that p^* for every decision problem that one encounters can be monumental. So being persistent in such a manner may pay off in the long run even though there might be cases in which one is practically certain that p^* even when one should not be so. According to Harsanyi, one accepts a proposition p^* when one assigns practical certainty to it persistently, and one recommends other people to do the same (p. 16).

However, Harsanyi's concept of acceptance is not that of belief—he acknowledges that one may accept a proposition in which one has low credence. Harsanyi recognises that being practically certain that p^* might lead to an increase in expected utility even if one's credence in p^* is actually quite low—for example, it may sometimes be expedient to accept that a certain scientific theory is true even if one's credence in the theory being literally true is quite low (pp. 16-17). Harsanyi's notion of acceptance cannot play the the role of being the doxastic component of knowledge.

7.4 Objections and Replies

We have seen how Restricted Belief and Restricted Rational Belief can resolve the dilemma. Let me now consider some objections to the theses.

7.4.1 The Threat of Conditional Bets

According to the theses, for us to believe or to rationally believe that p , we have to have a conditional credence of 1 in p given b , where b is the base proposition. One might then claim that even for the chance of a marginal return, we should be willing to risk life, limb, and soul on the conditional bet that p given b . But, the objection goes, such a risk is too huge, and we are typically reluctant or would not be rational to undertake it. Hence, there is

still too little that we believe or rationally believe.

I maintain that if we really believe that p , then there would indeed be situations in which we are willing to risk life, limb, and soul on a conditional bet that p given b .¹² If we believe that there is an apple on the table, for example, then there would be situations in which we would be willing to risk life, limb, and soul on a conditional bet that there is an apple on the table given the relevant base proposition. If we display unwillingness in such situations, then perhaps all we believe is the probabilistically-qualified claim that there is a high objective probability that there is an apple on the table. To see that this is not too unpalatable a bullet to bite, let's say that we presuppose that our senses are currently not deceiving us, that whatever currently appears to be the case is an accurate reflection of reality. And suppose that our credence that it *seems* to us that there is an apple on the table is 1—as Descartes points out, if it seems to us that so-and-so, we cannot be wrong in thinking that it seems to us that so-and-so. In his words, if 'I am now seeing light, hearing a noise, feeling heat', then 'I certainly *seem* to see, to hear, and to be warmed (Descartes, 1996, p. 19; his emphasis). Now note that the proposition that whatever currently appears to be the case is an accurate reflection of reality and the proposition that it seems to us that there is an apple on the table entail that there is an apple on the table. Our risk of losing the conditional bet thus seems to be no greater than our risk of losing a bet on a simple tautology.

One might still worry that there would be no situation in which it is rational to risk life, limb, and soul on a conditional bet that p given b , even if it is obvious that b entails p . Perhaps one might worry, as Descartes did, that an evil demon might have powers so great that it can trick us into believing that a certain proposition that is in fact false is obviously or self-evidently. For

¹²These would be situations in which, among other things, we are not averse to betting, we care only about making money, and the very act of betting does not make p false.

example, Christensen (2007) voices such a worry when he writes:

I don't think I can rationally be absolutely certain that no drug or demon could make it seem to me that I'm seeing clearly and distinctly when in fact I'm contemplating a falsity. And to the extent that I cannot absolutely preclude [the possibility that in believing that everything is self-identical, I'm believing a false claim due to a cognitive mistake], I fall short of rational absolute certainty [even in the claim that everything is self-identical]. (p. 19)

I agree with Christensen that we should not assign zero credence to the possibility of there being a drug or demon that can make something seem self-evidently true to us when it is in fact false. But it does not follow that when it is genuinely self-evident to us that everything is self-identical, we are not rational in being absolutely certain that everything is self-identical. Consider the following analogies. Suppose that when you are drunk, you can't tell the difference between drunkenness and sobriety, and that when you're insane, you can't tell the difference between sanity and insanity. It does not follow that you can't distinguish between drunkenness and sobriety *when you're clearly sober*, and between sanity and insanity *when you're clearly sane*.¹³ Now suppose that you have taken a drug that makes it seem self-evident to you that *p* is true, even if it is in fact false. This is a case in which having made a cognitive mistake, you are not able to distinguish between something that is genuinely self-evident, and something that merely seems to be self-evident. Such a case seems possible, but it does not follow that you can't distinguish between the two when something is genuinely self-evident to you (for instance, that everything is self-identical), and you've made no cognitive mistake.

¹³As Bernard Williams (1996) asks, 'why should the fact that when we are dreaming, we cannot tell that we are, imply that we cannot be sure we are awake when we are awake?' (p. xv).

One might point out that drunkenness and sobriety feel different from the inside. The difference might be the key to why it is rational for a sober person, but not for a drunk person, to be sure that she's in one state rather than the other. And likewise for sanity.¹⁴ Might there be a drug so powerful that after taking it, merely apparent self-evidence feels the same from the inside as genuine self-evidence feels to one who has not taken the drug? No. Suppose that there are propositions that are genuinely self-evident. Then there *must* be a real difference from the inside between merely apparent self-evidence and genuine self-evidence. Otherwise there will be no way to tell the two apart, and any knowledge that a purportedly self-evident proposition is true is contingent upon the assumption that it does not merely appear to be self-evident. But this is to say that it is not genuinely self-evident at all. Now one might just flat-out deny that there is anything self-evident. But there's no presumption in favour of such a denial. In fact, here is a presumption against it: it is self-evident that everything is self-identical.

There might still be a lingering intuition that it is never rational to bet life, limb, and soul on everything being self-identical or on obvious entailments. But such an intuition can be explained by the fact that bets are typically placed on declarative sentences or on assertive utterances. Someone might ask us, 'Do you want to place a bet on " p entails p "?' or 'Do you want to place a bet on "Everything is self-identical"?. We might be reluctant to accept such bets because we are unsure that we are really being offered a bet on p entailing p or on everything being self-identical. This is because we might be unsure what is really meant by ' p entails p ' or by 'Everything is self-identical'—an evil demon might trick us into thinking that such declarative sentences or assertive

¹⁴No doubt, it might not be rational for us to be *absolutely certain* that we're clearly sane even when we're so. But this is due more to the mundane reason that we might not be certain what constitutes sanity than to an evil demon making us unable to discriminate between a clear-cut state of sanity and a clear-cut state of insanity (whether or not we know whether those states are states of sanity or not).

utterances mean what they do not in fact mean.¹⁵

7.4.2 Still Too Few Beliefs and Rational Beliefs

Even if one is satisfied that the above objection is met, one might worry that the theses still yield too few beliefs or rational beliefs. Recall the discussion of the Preface Paradox in chapter 5. One might worry that to preserve closure in the face of Preface-like cases, we have to hold that non-fiction authors do not believe outright every claim they assert in their books. In general, the stronger the base proposition that *S* presupposes in a context of inquiry, the more difficult it will be to have a high rational credence in it. But the weaker the base proposition, the fewer beliefs or rational beliefs we have in that context. For example, if I presuppose the base proposition b_1 that my ticket is a loser, and your ticket is a loser, and his ticket is a loser, and so on, then I might well believe or be rational to believe that I will not be able to afford a round-the-world trip, that neither will you, that neither will he, and so on. But if I presuppose the base proposition b_2 that my ticket is a loser, then I might not believe or it will not be rational for me to believe that you will not be able to afford such a trip, that he will not be able to afford such a trip, and so on. For in such a case, my credence that you will not be able to afford such a trip conditional on b_2 , my credence that he will not be able to afford such a trip conditional on b_2 , and so on, are all less than 1.

Admittedly, in any context of inquiry, we may have fewer beliefs or rational beliefs than one may have thought. But first, within each context of inquiry, we will still have more beliefs and rational beliefs than we would have if belief and rational belief require full credence. And second, we will have many beliefs and rational beliefs across various contexts of inquiry. For a

¹⁵This is reminiscent of Stalnaker's (1984) metalinguistic approach to explaining why even though we believe all logical truths, it sometimes seems to us that we do not. See Stalnaker, 1984, pp. 84-87, and section 3.1 of chapter 3 in this dissertation.

proposition that you do not believe or are not rational to believe in a context of inquiry, there might well be another context of inquiry in which you do believe it. do not rationally believe. For example, in a context of inquiry in which I am asking whether I can afford a round-the-world trip, I may not be interested in the issue of whether you or he or she may afford such a trip. Suppose that in such a context, I do not presuppose that your ticket is a loser. Then I do not believe that you will not be able to afford a round-the-world trip. But there will be other contexts of inquiry in which I do, say, in a context in which I am primarily interested in the question of whether *you* will be able to afford such a trip.

One might still worry about having relatively few beliefs or rational beliefs within a single context of inquiry. But how deep does the worry run? Having few beliefs or rational beliefs does not mean that one is faced with doubt and scepticism, for there will still be many propositions in which one has high credence and is rational in having high credence. For example, consider a context of inquiry in which I am asking whether I can afford a round-the-world trip, and I have presupposed only that my ticket is a loser. I may still be rational in being very confident, even though I do not believe outright, that you will not be able to afford such a trip.

What's important is whether we have enough rational beliefs for us to be epistemically blameless in engaging in assertion and reasoning. And it seems that we do. If I am deliberating about what I should do next year, a rational belief that I cannot afford a round-the-world trip serves me well enough—I can be epistemically blameless in employing the corresponding premise in assertion and in reasoning, and have no need for beliefs about whether you or he or she can afford such a trip. One might worry that the lack of rational beliefs leads to a lack of knowledge. But why worry about a lack of knowledge in matters that have little bearing on one's present inquiries? Insofar as I remain in a context

in which I am deliberating about what to do for my vacation, knowledge about whether I can afford a round-the-world trip is what matters to me. Lack of knowledge that you can afford such a trip (due to an absence of the relevant rational belief) does not affect my inquiry. Of course, when I turn to inquiring about whether you can afford it, I may well be rational in presupposing that your lottery ticket is a loser. In such a case, I may well rationally believe and potentially know that you cannot afford the trip.

7.4.3 Moving the Bump Under the Rug?

One might think that my accounts of belief and of rational belief just move the bump under the carpet, and that a version of the dilemma that confronts attempts to relate belief and credence now arise at the level of presupposition. For instance, one might claim that just as we had to worry previously about how high a credence we need to have in a proposition for us to be rational to believe it, we now have to worry about how high a credence we need to have in a proposition for us to be rational to presuppose it. Also, one might claim that just as there was the challenge of maintaining that rational beliefs are closed under joint entailment without holding that rational beliefs require full credence, there is now a similar challenge of maintaining that rational presupposition is closed under joint entailment without holding that rational presupposition requires full credence.

However, it's not my aim to cash out presupposition in terms of credence or rational presupposition in terms of rational credence. I'm not claiming that to presuppose something is to assign it a high enough credence, and there's no dilemma for presupposition in the same way as there is one for belief. Recall also that according to Restricted Rational Belief, if it is rational for S to believe that p in a context of inquiry, it is rational for S to presuppose the conjunction of all relevant propositions and to assign it a credence close to 1. The thesis

in itself is neutral on whether rational presupposition is closed under joint entailment—for example, for all the thesis says, S may rationally presuppose that $p \wedge q \wedge r$ and that s , but not that $p \wedge q \wedge r \wedge s$. But if $p \wedge q \wedge r$ is the conjunction of all relevant propositions, and S assigns high enough credence to it, S may well rationally believe any proposition in which her rational credence conditional upon $p \wedge q \wedge r$ is 1.

7.4.4 Animals and Beliefs

One might object that the theses do not allow most animals to count as having binary beliefs, since they tie such beliefs so closely to human activities like deliberation and assertion from which, arguably, most animals are excluded. But remember that I am interested in a notion of rational belief that fulfils certain epistemological roles, bridging rational credence with knowledge and with epistemically blameless assertion and reasoning. Given this interest, I am happy to attribute credences and desires to animals, but since they do not engage in reasoning and assertion, and arguably do not possess propositional knowledge (especially on a view of justification that requires that we believe that we have justification for p to know that p), there is no point in talking about what they believe or are rational to believe.

In adopting such a stance, I am accompanied by several partners in crime. Ronald B. de Sousa (1971), who equates a belief that p with a disposition to assent that p , also seems happy to say that animals have credences and desires but not binary beliefs (pp. 57-58; p. 64). In chapter 4, I examined some attempts to carve out a functional role for beliefs by tying them to reasoning and to assertion. Williamson (2000), Wedgwood (2008), Smithies (forthcoming), and Holton (2008) all hold that if one believes that p , then one is disposed to use p as a premise in reasoning, whereas if one has a merely high credence that p , one will not be disposed to do so outright. If what they

all mean by reasoning or deliberation is some sort of deliberative, reflective reasoning, then on their views, most animals do not have beliefs, assuming that most animals do not engage in such reasoning. We have also seen that Kaplan (1996) tries to tie belief very closely to assertion. But if animals are not capable of assertion, then on his view, it is hard to make sense of how they may have beliefs.

One might think that if animals possess doxastic states at all, they possess them in the form of beliefs rather than credences: since credences are more fine-grained than beliefs, possessing the former would be harder than possessing the latter. However, even though the notion of credence is arguably more sophisticated than the notion of belief, it doesn't follow that possessing something picked out by a more sophisticated notion is harder than possessing something picked out by a less sophisticated one. Also, animals may act according to their credences and desires without calculating or deliberating with such credences and desires consciously. Perhaps there is an additional worry that there are just too many credences to be stored in an animal's head, but I've dealt with such a worry in section 4.3.1 of chapter 4.

Now I do not wish to deny that we sometimes appeal to binary beliefs to give a coarse-grained account of the action and behaviour of animals. Perhaps we sometimes pretend that they are creatures capable of deliberate and reflective reasoning. We then project our own reasoning processes onto them, and use such projection as a way by which to explain and predict how they would act. I am also willing to grant that people in daily life may sometimes use 'belief' to mean a state of high confidence. If a dog desires to chew on a bone, and has high confidence that there is a bone buried at a spot next to the plum tree in the garden, we might well explain why the dog is digging at the spot by saying that it believes that there is a bone buried underneath. But as mentioned in chapter 5, I'm more interested in a notion of belief that

obeys Belief-Reason and Belief-Assert than in what ‘belief’ means in ordinary language.

7.4.5 Retraction Worries and Mixed Situations

According to the standard contextualist treatment of *knowledge* ascriptions, ‘S knows that *p*’ may be true in a low-standards context but false in a high-standards context. But John MacFarlane (2005) claims that if S asserted ‘I know that *p*’ earlier, and standards have now been raised so that she no longer knows that *p*, we would expect her to *retract* her earlier assertion and concede that ‘I know that *p*’, as asserted by her earlier, was false (pp. 202-203). This phenomenon of retraction, MacFarlane holds, ‘forms the basis of standard arguments against contextualism’ (p. 203). For according to contextualism, there should be no need for S to retract her assertion—even though she does not know that *p*, and did not know that *p* earlier, ‘I know that *p*’, as asserted by S in the earlier context, remains true (*ibid.*).

One might worry that Restricted Belief and Restricted Rational Belief face a rather similar problem, since according to the theses, ‘S believes that *p*’ and ‘S rationally believes that *p*’ may be true in some contexts of inquiry but not in others.¹⁶ But there is no problem, since there is no similar phenomenon of retraction. Suppose, for example, that while not thinking about the no-loss bet at time t_1 , Jack is correct in asserting both ‘I believe there’s flour in the cupboard’ and ‘I rationally believe there’s flour in the cupboard’. Suppose that at t_2 , he enters into a context of inquiry in which he entertains the no-loss bet. In such a context, he does not believe, and hence does not rationally believe, that there’s flour in the cupboard. But he need not retract his earlier assertions. When Jack considers the no-loss bet, it’s natural to describe him as

¹⁶I’m not claiming that MacFarlane’s objection to contextualism is fatal. But it is at least a *prima facie* problem for contextualism. The worry is whether there is a similar *prima facie* problem for the theses I defend.

losing his belief that there's flour in the cupboard—and he can't *lose* something he does not have in the first place.

One might grant the preceding and still think that Jack should retract his earlier assertion that he was *rational* in having the belief. But he shouldn't, since even at t_2 , his assertion at t_1 was perfectly in order. Recall that as I am using the term 'rational', to be rational in believing that p is to be epistemically blameless in believing that p . Now consider the case of moral blameworthiness as an analogy. If at t_1 , Jack is morally blameless in performing a certain act, he may not be morally blameless in performing the same act at some other time and in some other kind of situation. But at any other time and in any other kind of situation, it remains that he's morally blameless in performing that act at t_1 . Analogously, if at t_1 , Jack is epistemically blameless in believing that p , he may not be epistemically blameless in doing so at some other time and in some other kind of situation. But at any other time and in any other kind of situation, it remains that he's epistemically blameless in believing that p at t_1 .

According to what MacFarlane calls *sensitive invariantism*, whether S knows that p depends on the situation S finds herself in (p. 202). But MacFarlane claims that we do not say things like 'John knows that he won't be able to afford health insurance, but if he were discussing the possibility that he might win the lottery, he would not know this [while still retaining his belief that he won't be able to afford health insurance]' (p. 202). According to MacFarlane, such observations 'form the basis of standard arguments against sensitive invariantism' (*ibid.*).

One might worry that Restricted Belief and Restricted Rational Belief face a rather similar problem, since according to the theses, whether 'S believes that p ' or 'S rationally believes that p ' is true depends on S's context of inquiry rather than on the belief attributor's. But once again, there is no

problem. According to Restricted Belief, we may say things like ‘John believes that he won’t be able to afford health insurance, but if he were discussing the possibility that he might win the lottery, he would not believe this’, and according to Restricted Rational Belief, we may say things like ‘It is rational for John to believe that he won’t be able to afford health insurance, but if he were discussing the possibility that he might win the lottery, it would not be rational for him to believe this’. The first sentence sounds perfectly fine. There is nothing wrong with holding that S believes that p in the current situation, but would not believe that p in some other situation. I don’t have firm intuitions about the second sentence. But again, recall that as I am using the term ‘rational’, to be rational in believing that p is to be epistemically blameless in believing that p . Do we say things like ‘It is epistemically blameless for John to believe that he won’t be able to afford health insurance, but if he were discussing the possibility that he might win the lottery, it would not be epistemically blameless for him to believe this’? Well, we do, but usually when we are doing philosophy—ordinary speakers hardly use terms such as ‘epistemically blameless’. The issue of whether it is fine to say such things should be guided by our best theories of belief, and not merely on ordinary linguistic intuitions (which seem murky in this case).¹⁷

¹⁷In a similar vein, Wedgwood (2008) writes:

The term ‘justified belief’ is simply not in such common use as the term ‘knowledge’. Unlike the terms ‘knowledge’ or ‘belief’, the term ‘justified belief’ is not heard every day on practically everyone’s lips. It is primarily *philosophers* who use the term. Moreover, philosophers have used the term to express many different notions. For this reason, when philosophers use the term, they have to clarify exactly which of these many different notions they are using the term to express. Because the term is capable of expressing so many different notions, and is anyway chiefly used by philosophers rather than by ordinary speakers, it is not surprising that our intuitions relating to the them are less clear-cut than our intuitions relating to the term ‘know’. (p. 7)

7.4.6 Beliefs and Context-Relativity

Some philosophers do not think much of the view that binary beliefs are context-sensitive. For example, Maher (1986) thinks that ‘belief seems clearly *not* to be a context-relative notion’ (p. 383; emphasis his).¹⁸ Also, Kaplan (1996) considers a version of Restricted Belief-Act that is context-relative, but dismisses it on the grounds that instead of accounting for a person’s ‘believing *p*, *period*’, it accounts for some other notion, such as that of a person ‘believing *p* in a given context’ (p. 105; emphasis his). Similarly, Frankish (2009) claims that binary belief, ‘as commonly understood, is unqualified as to context’ (p. 9). And Michael Bratman (1992) writes, ‘Reasonable belief is, in an important way, context independent: at any one time a reasonable agent normally either believes something (to degree *n*) or does not believe it (to that degree). She does not at the same time believe that *p* relative to one context but not relative to another’ (p. 3).

Although the above-mentioned philosophers have the intuition that belief is not context-relative, the dilemma discussed in chapter 5 might be reason to question whether the intuition is worth preserving. Or at least, the best hope of dodging both horns of the dilemma might be to hold that there is an important notion of belief, namely, the epistemic notion, that is context-relative. Indeed, as we have seen earlier on, many accounts of how (rational) belief and (rational) credence are related hold that (rational) belief is context-relative in some way or another.¹⁹

Admittedly, Bratman articulates a fair worry about belief being context-relative when he points out that a reasonable agent does not at the same time

¹⁸Maher now seems to have become sympathetic to the view that belief is sensitive to context. In particular, he seems to endorse the Contextualised Lockean Thesis. See Maher (2005), pp. 6-7.

¹⁹Someone like Harman (1986) might take the dilemma to be reason to understand credences as beliefs about objective probabilities. I’ve argued in chapter 2 why such a view should be rejected.

believe that p relative to one context but not relative to another. On such a view of belief, the very same person, in the very circumstances that she is in, may well be correctly judged by Jack on Monday to believe that p , and correctly judged by Ida on Tuesday not to believe that p . We cannot say, as Kaplan puts it, whether this person believes that p , period, or does not believe that p , period.

Now earlier in this chapter, I argued against Wedgwood's account of belief, which allows beliefs to be context-relative in such a way. But the accounts of belief defended by Weatherson (2005), by Kyburg (1998), and by me, among many others, are not context-relative in such a sense. According to such views, whether a person believes that p at a particular time depends on the situation she finds herself in at that time. If the situation had been different (but her credence in p and her evidence for p remain the same), she might not have believed that p . But this is not to say that at the same time, she both believes that p , relative to her current situation, and also does not believe that p , relative to another.

Bratman grants that a person's beliefs can change when the context she is in changes, but in a sense that does not make the beliefs context-relative in any interesting sense. He writes:

Richmond Thomason (1986, p. 345) discusses an example of someone who, when he thinks about dentist appointments, believes he has a dentist appointment on Friday, but when he thinks about Friday does not believe this. But such changes will likely get this person in trouble in his planning, and are not a case of reasonably believing that p in one context but not in another. [...] Or again, perhaps I earlier believed that I had a copy of the book at home and now no longer believe it. But this is a change of mind; it is not that I now believe this relative to one context but not relative

to another. (p. 3).

I agree that the examples Bratman offers are not examples of a context-relative notion of binary belief that is philosophically interesting. To get such a notion, we want cases in which changes in binary belief will not get a person into trouble in her planning and which do not involve any real change of mind. But how are such cases possible?

According to Bratman, the view that belief is not context-relative is ‘part and parcel of our ordinary conception of belief—be it all-or-none or a matter of degree’ (p. 4). I maintain that beliefs are context-relative in at least one interesting sense, but agree with what Bratman says about graded beliefs. But what is involved in change of mind and in planning are really credences and not binary beliefs (given that credences do all the heavy lifting). Now suppose we have a case in which one has a binary belief that p in one context, but not in another, even though one’s credences and desires remain the same in both contexts. Since one’s credence in p remains the same, there is in an important sense no real change of mind. Neither will one get into trouble in planning, if one acts according to credences and desires that are reasonable. So this is just the kind of case we want. And we have witnessed such cases—recall the example involving Jack’s belief about whether there is flour in the cupboard.

It is also not clear that Kaplan’s (1996) account of belief precludes thinking of belief as context-relative (in the sense that Weatherson’s (2005) and Kyburg’s accounts of belief, for example, are context-relative). According to him, what one believes depends on one’s preferences regarding what to assert. Keeping the amount of evidence that one has for the truth of p fixed, whether a person believes p will depend on how she balances her desire to avoid error and her desire to be as comprehensive as possible.²⁰ But how she

²⁰Recall the discussion of Kaplan’s account of belief in chapter 3.

tries to achieve this balance may well depend on the context of her inquiry—sometimes she might give more weight to being comprehensive, and sometimes she might give more weight to being accurate.

Finally, it is not clear that Frankish's own account of belief allows him to avoid claiming that belief is context-relative. He thinks that believing that p involves adopting a premising policy with respect to p , i.e. a policy of employing p as a premise in one's deliberations.²¹ In particular, for one to believe that p , one has to be willing to rely on p in what he calls *TCP deliberations*, where a TCP (truth-critical with respect to premise) deliberation is

one where we attach an overridingly high desirability to taking only truths as premises—for example, a situation where we are required to predict the effects of a certain action and where there are large penalties for false predictions (p. 14).

But there will be TCP deliberations in which one is willing to rely on p as a premise, and others in which one is not so willing, even though one's evidence about whether p is true remains the same. For example, suppose it is really important that Jack has enough flour in the cupboard to make a pizza. (He wants to propose to his partner, who loves pizza.) His credence that there is flour in the cupboard might be high enough for him to rely on the premise that there is flour in the cupboard in his TCP deliberation regarding whether to drop by at the supermarket or to go straight home. But presumably, he is not willing to use the same premise in a TCP deliberation regarding whether to accept or reject the no-loss bet.

Frankish recognises that just because one believes that p does not mean that one is willing to use p as a premise in any TCP deliberation whatsoever. He admits that '[t]here will be situations where there is little to gain from

²¹Cf. Harsanyi (1985).

acting on a certain premise or goal and much to lose' (p. 17). But he thinks that 'it is plausible to regard premising policies as having built-in exception clauses, excluding those deliberations where it would obviously be dangerous or absurd to rely on the premise or goal in question'. But notice that in deliberating whether to accept or reject the no-loss bet, it is *not* dangerous to employ the premise that there is flour in the cupboard, since the stakes are not that high. And if it is absurd for Jack to reject the bet based on the premise that there is no flour in the cupboard, it is precisely because in the context of considering the bet, it is not rational for him to believe the premise outright. Appealing to 'built-in exception clauses' seems *ad hoc* at best.

Chapter 8

Conclusion

We often appeal to our beliefs and desires to explain or predict the way we act and behave. In such explanation or prediction, it is natural to think that beliefs and desires come in various strengths. I've argued that, in fact, the *only* kind of belief we need to appeal to is the kind that comes in various strengths, namely, credence.

In chapter 2, we saw that there is work that credences do that binary beliefs can't do. Views according to which credences are just binary beliefs about objective probabilities are problematic. There are cases in which one's credence in a proposition is x , for some $x \in [0, 1]$, but one does not believe outright that the objective probability of p being true is x .

In chapter 3, I addressed the worry that we can't appeal to credences or rational credences to explain the action and behaviour of non-ideal agents like human beings. One might think that in using subjective probabilities to model the credences of a rational agent, we impute logical omniscience or an unrealistic level of logical sophistication to the agent. But it is not compulsory that we model credences using the standard probability axioms. In that chapter, I proposed more realistic constraints on credences than those given by the standard axioms. I also argued that my constraints' lack of mathematical structure (compared to the constraints provided by the axioms) is not a problem, but an advantage.

Given that credences are more fine-grained than binary beliefs, it's natural to think that we do not need to appeal to binary beliefs over and above an appeal to credences to explain and predict action and behaviour. However, as we saw in chapter 4, some claim that credences are conceptually dependent upon binary beliefs. Some of the arguments that are supposed to support this claim rely on the thought that without binary beliefs, there will be too many credences for our brains to store and manipulate. I explained why such a thought is unwarranted. Some philosophers of mind hold that beliefs are not stored in the brain. Some hold that they are stored in a map-like rather than sentence-like way—and even a simple map can encode a vast amount of information. Furthermore, even if *explicit* credences were stored in the brain in a sentence-like way, there would still be a vast number of *implicit* credences that are not stored in such a manner.

Some also claim that binary belief has a non-redundant functional role to play in deliberation and in assertion. But I argued that since deliberative reasoning and assertion are *acts* in which we engage, our engagement in them can be accounted for without recourse to talk about binary belief—we perform such acts because doing so leads to greater expected utility than not doing so.

Even though one may be satisfied that credences rather than binary beliefs do all the work in explaining and predicting action and behaviour, one may still think that *pace* Jeffrey, not all the marrow has been sucked out of the notion of binary belief. For one may think that beliefs have a role to play in epistemology. After all, knowledge that p is thought to entail a rational *belief* that p , and it is plausible that if one rationally believes that p , then it is epistemically blameless for one to assert that p or to use p as a premiss in reasoning. To investigate how much marrow is left in the notion of binary belief, I investigated various accounts of how belief is related to credence—if we can cash out binary belief in terms of credences and other

notions, its role in epistemology will be diminished. In chapter 5, I considered the Threshold View, according to which belief is just sufficiently high credence, and the Lockean Thesis, according to which rational belief is just sufficiently high rational credence. I argued that they succumb to a dilemma that also poses a challenge to anyone who holds that to believe or to rationally believe that p , it is *necessary* that one's credence in p meets a high enough threshold. If belief or rational belief requires full credence, then we have very few beliefs or rational beliefs. But if not, then among other things, one is faced with the lottery paradox, as well as the failure of certain closure principles.

In chapter 6, I considered various other theories of the relationship between belief and credence. According to some theories, to believe that p is to have full credence in p , but such credence is defined relative to a limited set of possibilities. I argued that this leaves the theories unable to account for certain behavioural dispositions that are best explained by holding that some credence is assigned to possibilities outside the limited set. Other theories, such as the Contextualised Threshold View and the Contextualised Lockean Thesis are unable to preserve closure. Yet other theories, such as Belief-Act and Belief-Preference, get impaled on the first horn of the dilemma by requiring that belief or rational belief be accompanied by full credence; I also argued that although restricted versions of the theories avoid the first horn of the dilemma, they get the wrong results in certain cases, delivering the verdict that an agent believes or rationally believes something that intuitively she doesn't.

Finally in chapter 7, I defended my own account of the relationship between belief and credence, arguing that it resolves the dilemma. According to the account, to (rationally) believe that p is to have a (rational) credence of 1 in p , conditional on some special proposition that I called the base proposition. What this base proposition is depends on the context of inquiry in which an agent finds herself. I also anticipated and answered some objections to my

account. If the account is correct, then instead of talking about beliefs, we can talk about credences in various contexts of inquiry. This in turn means that even in epistemology, binary beliefs play a less important role than one might have thought.

Bibliography

- Adams, Ernest W. (1998). *A Primer of Probability Logic*. Center for the Study of Language and Information.
- Alston, William (2005). *Beyond "Justification": Dimensions Of Epistemic Evaluation*. Ithaca, New York: Cornell University Press.
- Arló-Costa, Horacio (2001). 'Bayesian Epistemology and Epistemic Conditionals: On the Status of the Export-Import Laws'. *The Journal of Philosophy*, 98: 11, 555-593.
- Bartha, Paul (2004). 'Countable Additivity and the de Finetti Paradox'. *British Journal for the Philosophy of Science*, 55, 301-321.
- Beall, JC & Restall, Greg (2005). 'Logical Consequence'. In *The Stanford Encyclopedia of Philosophy (Winter 2005 Edition)*. Ed. Edward N. Zalta. URL = <http://plato.stanford.edu/archives/win2005/entries/logical-consequence/>.
- Braddon-Mitchell, David and Frank Jackson (2007). *Philosophy of Mind and Cognition: An Introduction*. UK: Blackwell Publishing.
- Braithwaite, Richard B. (1946). 'Belief and Action'. *Aristotelian Society*, supplementary volume 20, 1-19.
- Bratman, Michael E. (1992) 'Practical Reasoning and Acceptance in a Context'. *Mind*, 101, 1-15.
- Christensen, David (2004). *Putting Logic in its Place: Formal Constraints on Rational Belief*. New York: Oxford University Press.
- Carnap, Rudolf (1950). *Logical Foundations of Probability*. Chicago: Univer-

sity of Chicago Press.

- Cherniak, Christopher (1981). 'Minimal Rationality'. *Mind*, 90, 161-183.
- Christensen, David (2004). *Putting Logic In Its Place*. New York: Oxford University Press.
- Christensen, David (2007). 'Does Murphy's Law Apply in Epistemology? Self-Doubt and Rational Ideals'. In *Oxford Studies in Epistemology Volume 2*. Ed. Tamar Szabó Gendler and John Hawthorne. New York: Oxford Clarendon Press. 3-31.
- Churchland, Paul M. (1981). 'Eliminative Materialism and the Propositional Attitudes'. *The Journal of Philosophy*, 78:2, 67-90.
- Davidson, Donald (1985). 'A New Basis for Decision Theory'. *Theory and Decision*, 18, 87-98.
- De Finetti, B. (1937/1980). 'Foresight. Its Logical Laws, Its Subjective Sources'. In *Studies in Subjective Probability*. Ed. H. E. Kyburg, Jr. and H. E. Smokler. Robert E. Krieger Publishing Company. 93-158.
- DeRose, Keith (1996). 'Knowledge, Assertion and Lotteries'. *Australasian Journal of Philosophy*, 74: 4, 568-580.
- DeRose, Keith (2002). 'Assertion, Knowledge, and Context'. *The Philosophical Review*, 111:2, 167-203.
- deSousa, Ronald B. (1971). 'How to Give a Piece of Your Mind: Or the Logic of Belief and Assent'. *Review of Metaphysics*, 25:1, 52-79.
- Dennett, Daniel. (1978). *Brainstorms: Philosophical Essays on Mind and Psychology*. Cambridge, MA: MIT Press.
- Dennett, Daniel (1991). 'Real patterns'. *The Journal of Philosophy*, 88:1, 27-51.
- Descartes, Rene (1996). *Meditations on First Philosophy: With Selections from the Objections and Replies*. Ed. John Cottingham. UK: Cambridge University Press.

- Douven, Igor (2002). 'A New Solution to the Paradoxes of Rational Acceptability'. *British Journal for the Philosophy of Science*, 53, 391-410.
- Eells, Ellery (1990). 'Bayesian Problems of Old Evidence'. In *Minnesota Studies in the Philosophy of Science*, 14. Ed. C. Wade Savage. 205-223.
- Evine, Simon J, (1999). 'Believing Conjunctions'. *Synthese*, 118, 201-227.
- Fantl, Jeremy and Matthew McGrath (2002). 'Evidence, Pragmatics, and Justification'. *The Philosophical Review*, 111:1, 67-94.
- Field, Hartry. (2009). 'Pluralism in Logic'. *The Review of Symbolic Logic*, 2:2, 342-359.
- Fishburn, Peter C. (1981). 'Subjective Expected Utilities: A Review of Normative Theories'. *Theory and Decision*, 13, 139-199.
- Fodor, Jerry (1989). *Psychosemantics*. Cambridge, MA: MIT Press.
- Foley, Richard (1993). *Working Without a Net: A Study of Egocentric Epistemology*. New York: Oxford University Press.
- Frankish, Keith (2009). 'Partial Belief and Flat-Out Belief'. In *Degrees of Belief*. Ed. F. Huber and C. Schmidt-Petri. Dordrecht: Springer. 75-93.
- Gaifman, Haim (2004). 'Reasoning with Limited Resources and Assigning Probabilities to Arithmetical Statements'. *Synthese*, 140, 97-119.
- Ganson, Dorit (2008). 'Evidentialism and Pragmatic Constraints on Outright Belief'. *Philosophical Studies*, 139, 441-458.
- Garber, Daniel (1983). 'Old Evidence and Logical Omniscience in Bayesian Confirmation Theory'. In *Minnesota Studies in Philosophy of Science*, 10. 99-131.
- Good, I.J. (1986). *Good Thinking: The Foundations of Probability and its Applications*. Minneapolis, MN: University of Minnesota Press.
- Hacking, Ian (1967). 'Slightly More Realistic Personal Probability'. *Philosophy*

of Science, 34:4, 311-325.

Hájek, Alan (1997). “‘Mises Redux’—Redux. Fifteen Arguments Against Finite Frequentism’, *Erkenntnis*, 45, 209-227.

Hájek, Alan (2001). ‘Probability, Logic, and Probability Logic’. In *The Blackwell Companion to Logic*. Ed. Lou Goble. Blackwell, 362-384.

Hájek, Alan (2003). ‘What Conditional Probability Could Not Be’. *Synthese*, 137, 273-323.

Hájek, Alan (2006). Interview in *Masses of Formal Philosophy*. Ed. Vincent F. Hendricks and John Symons. Automatic Press.

Hájek, Alan (2009). ‘Interpretations of Probability’. In *The Stanford Encyclopedia of Philosophy (Spring 2009 Edition)*. Ed. Edward N. Zalta. URL = <http://plato.stanford.edu/archives/spr2009/entries/probability-interpret/>.

Hájek, Alan and Michael Smithson (forthcoming). ‘Rationality and Indeterminate Probabilities’.

Hájek, Alan. ‘A Puzzle About Belief’. Manuscript.

Harman, Gilbert (1986). *Change in View: Principles of Reasoning*. Cambridge, MA: MIT Press.

Harman, Gilbert (1970). ‘Language Learning’. *Nous*, 4:1, 33-43.

Harsanyi, John C. (1985). ‘Acceptance of Empirical Statements: A Bayesian Theory Without Cognitive Utilities’. *Theory and Decision*, 18, 1-30.

Hawthorne, James (2009). ‘The Lockean Thesis and the Logic of Belief’. In *Degrees of Belief*. Ed. F. Huber and C. Schmidt-Petri. Dordrecht: Springer. 49-74.

Hawthorne, James and Luc Bovens (1999). ‘The Preface, the Lottery, and the Logic of Belief’. *Mind*, 108, 241-264.

Hawthorne, John (2004) *Knowledge and Lotteries*. Oxford: Oxford University Press.

- Hawthorne, John and Brian Weatherson. 'Beliefs Old and New'. Manuscript.
- Hill, Christopher S. and Joshua Schechter (2007). 'Hawthorne's Lottery Paradox and the Nature of Belief'. *Philosophical Issues*, 17, 102-122.
- Hintikka, Jaakko (1975). 'Impossible Possible Worlds Vindicated'. *Journal of Philosophical Logic*, 4, 475-484.
- Holton, Richard (2008). 'Partial Belief, Partial Intention'. *Mind*, 117, 27-58.
- Hunter, Daniel (1996) 'On the Relation Between Categorical and Probabilistic Belief'. *Nous*, 30: 1, 75-98.
- James, William (1956). 'The Will to Believe: and Other Essays in Popular Philosophy, and Human Immortality'. New York: Dover Publications.
- Jeffrey, Richard. (1970). 'Dracula Meets Wolfman: Acceptance vs. Partial Belief'. In *Induction, Acceptance, and Rational Belief*. Ed. Marshall Swain. Dordrecht, Holland: D. Reidel Publishing Company. 157-185.
- Jeffrey, Richard. (1983). *The Logic of Decision*. Chicago: University of Chicago Press.
- Jeffrey, Richard. (1992). 'Radical Probabilism (Prospectus for a User's Manual)'. *Philosophical Issues*, 2, 193-204.
- Jeffrey, Richard. (1997). 'Unknown Probabilities'. *Erkenntnis*, 45, 327-335.
- Kaplan, Mark (1996). *Decision Theory as Philosophy*. New York: Cambridge University Press.
- Kitcher, Philip and Achille Varzi (2000). 'Some Pictures are Worth 2^{\aleph_0} Sentences'. *Philosophy*, 75, 377-381.
- Kyburg Jr., Henry E. (1970). 'Conjunctivitis'. In *Induction, Acceptance, and Rational Belief*. Ed. Marshall Swain. Dordrecht, Holland: D. Reidel Publishing Company. 55-82.
- Kyburg Jr., Henry E. (1988). 'Full Belief'. *Theory and Decision*, 25, 137-162.
- Lackey, Jennifer (2007). 'Norms of Assertion'. *Nous*, 41:4, 594-626.
- Lance, Mark Norris (1995). 'Subjective Probability and Acceptance'. *Philo-*

sophical Studies, 77, 147-179.

Lasonen-Aarnio, Maria (2008). 'Single Premise Deduction and Risk'. *Philosophical Studies*, 141, 157-173.

Levi, Isaac (1980). *The Enterprise of Knowledge: An Essay on Knowledge, Credal Probability and Chance*. Cambridge, MA: MIT Press.

Levi, Isaac (1985). 'Imprecision and Indeterminacy in Probability Judgment'. *Philosophy of Science*, 52:3, 390-409.

Levi, Isaac and Sidney Morgenbesser (1964). 'Belief and Disposition'. *American Philosophical Quarterly*, 1:3, 221-232.

Lewis, David (1973). *Counterfactuals*. UK: Oxford, Basil Blackwell.

Lewis, David (1974). 'Radical Interpretation'. *Synthese*, 23, 331-344.

Lewis, David (1979). 'Scorekeeping in a Language Game'. *Journal of Philosophical Logic*, 8, 339-59.

Lewis, David (1980). 'A Subjectivist's Guide to Objective Chance'. In *Studies in Inductive Logic and Probability*, 2. Ed. Richard C. Jeffrey. Berkeley and Los Angeles: University of California Press. 263-293.

Lewis, David (1986). *On the Plurality of Worlds*. Oxford: Blackwell.

Lewis, David (1994). 'Reduction of Mind'. In *A Companion to the Philosophy of Mind*. Ed. S. Guttenplan. Oxford: Blackwell. 412-429.

Lewis, David (1996). 'Elusive Knowledge'. *Australasian Journal of Philosophy*, 74:4, 549-567.

Lewis, David (1998). 'Relevant Implication'. In *Papers in Philosophical Logic*. Cambridge University Press. 111-124. First published in *Theoria*, 54, (1988), 161-174.

MacFarlane, John (2005). 'The Assessment-Sensitivity of Knowledge Attributions'. In *Oxford Studies in Epistemology Volume 1*. Ed. Tamar Szabó Gendler and John Hawthorne. New York: Oxford Clarendon Press. 197-233.

- Maher, Patrick (1986). 'The Irrelevance of Belief to Rational Action'. *Erkenntnis*, 24, 363-384.
- Maher, Patrick (1993). *Betting on Theories*. New York: Cambridge University Press.
- Maher, Patrick (2006). 'Review of *Putting Logic In Its Place*, by David Christensen'. *Notre Dame Journal of Formal Logic*, 47:1, 133-149.
- Makinson, David C. (1965). 'The Paradox of the Preface'. *Analysis*, 25, 205-207.
- Marcus, Ruth Barcan (1983). 'Rationality and Believing the Impossible'. *The Journal of Philosophy*, 80:6, 321-338.
- McGee, Vann (1985). 'A Counterexample to Modus Ponens'. *The Journal of Philosophy*, 82:9, 462-471.
- Nelkin, Dana K. (2000). 'The Lottery Paradox, Knowledge, and Rationality'. *The Philosophical Review*, 109:3, 373-409.
- Nolan, Daniel (1997). 'Impossible Worlds: A Modest Approach'. *Notre Dame Journal of Formal Logic*, 38:4, 535-572.
- Nover, Harris and Alan Hájek (2004). 'Vexing Expectations'. *Mind*, 113, 237-249.
- Nozick, Robert (1993). *The Nature of Rationality*. Princeton, NJ: Princeton University Press.
- Pollock, John L. (2006). *Thinking About Acting: Logical Foundations for Rational Decision Making*. New York: Oxford University Press.
- Pollock, John L. (forthcoming). 'Problems for Bayesian Epistemology'.
- Priest, Graham (2006). *In Contradiction: a Study of the Transconsistent* 2nd Ed. New York: Oxford University Press.
- Pryor, James (2001). 'Highlights of Recent Epistemology'. *British Journal for the Philosophy of Science*, 52, 95-124.
- Putnam, Hilary (1973). 'Reductionism and the Nature of Psychology'. *Cogni-*

tion, 2:1, 131-146.

Radford, Colin (1966). 'Knowledge—By Examples'. *Analysis*, 27, 1-11.

Ramsey, Frank P. (1926/1990). 'Truth and Probability'. In *Philosophical Papers*. Ed. D.H. Mellor. New York: Cambridge University Press. 52-96.

Roorda, Jonathan (1995). Manuscript. 'Revenge of Wolfman: A Probabilistic Explication of Full Belief'. URL = <http://www.princeton.edu/~bayesway/put/Wolfman.pdf>. Last accessed: 16th Dec. 2009.

Ryan, Sharon (1991). 'The Preface Paradox'. *Philosophical Studies*, 64, 293-307.

Ryan, Sharon (1996). 'The Epistemic Virtues of Consistency'. *Synthese*, 109, 121-141.

Sargent, Mark (2009). 'Answering the Bayesian Challenge'. *Erkenntnis*, 70, 237-252.

Schwarz, Wolfgang (2007). 'Beliefs and Thresholds' (Blog post). URL = <http://www.umsu.de/wo/2007/492>. Last accessed: 16th Dec. 2009.

Schwitzgebel, Eric (2002). 'A Phenomenal, Dispositional Account of Belief'. *Nous*, 36:2, 249-275.

Schwitzgebel, Eric (2008). 'Belief'. *The Stanford Encyclopedia of Philosophy* (Fall 2008 Edition). Ed. Edward N. Zalta. URL = <http://plato.stanford.edu/archives>

Skyrms, Brian (1977). 'Resiliency, Propensities, and Causal Necessity'. *The Journal of Philosophy*, 74, 704-713.

Smithies, Declan (forthcoming). 'The Normative Role of Knowledge'.

Stalnaker, Robert (1984). *Inquiry*. Cambridge: MA: The MIT Press.

Stalnaker, Robert (2002). 'Common Ground'. *Linguistics and Philosophy*, 25, 701-721.

Stalnaker, Robert (2003). *Ways a World Might Be: Metaphysical and Anti-Metaphysical Essays*. New York: Oxford University Press.

- Stanley, Jason (2008). 'Knowledge and Certainty'. *Philosophical Issues*, 18, 35-57.
- Sturgeon, Scott (2008). 'Reason and the Grain of Belief'. *Nous*, 42:1, 139-165.
- Talbott, William (2008). 'Bayesian Epistemology'. In *The Stanford Encyclopedia of Philosophy (Fall 2008 Edition)*. Ed. Edward N. Zalta. URL = <http://plato.stanford.edu/archives/fall2008/entries/epistemology-bayesian/>.
- van Fraassen, Bas (1989). *Laws and Symmetry*. Oxford: Clarendon Press.
- van Fraassen, Bas (1995). Fine-Grained Opinion, Probability, and the Logic of Full Belief. *Journal of Philosophical Logic*, 24, 349-377.
- Vogel, Jonathan (1990). 'Are There Counterexamples to the Closure Principle?'. In *Doubting*. Ed. M.D. Ross and G. Ross. Kluwer Academic Publishers. 13-27.
- Wallace, R. Jay (2009). 'Practical Reason'. *The Stanford Encyclopedia of Philosophy (Summer 2009 Edition)*. Ed. Edward N. Zalta. URL = <http://plato.stanford.edu/archives/sum2009/entries/practical-reason/>.
- Walley, Peter (1991). *Statistical Reasoning with Imprecise Probabilities*. UK: St. Edmundsbury Press.
- Weatherson, Brian (2003). 'From Classical to Intuitionistic Probability'. In *Notre Dame Journal of Formal Logic*, 44:2, 111-23.
- Weatherson, Brian (2005). 'Can We Do Without Pragmatic Encroachment?'. *Philosophical Perspectives*, 19, 417-443.
- Wedgwood, Ralph (2008). 'Contextualism About Justified Belief'. *Philosophers' Imprint*, 8:9, 1-20. URL = www.philosophersimprint.org/008009
- Weiner, Matthew (2005). 'Must We Know What We Say?'. *The Philosophical Review*, 114:2, 227-251.
- Weintraub, Ruth (2001). 'The Lottery: A Paradox Regained and Resolved'.

Synthese, 129, 439-449.

Williams, Bernard (1996). 'Introductory Essay'. In *Meditations on First Philosophy: With Selections from the Objections and Replies*. Ed. John Cottingham. UK: Cambridge University Press. vii-xvii.

Williamson, Timothy (2000). *Knowledge and Its Limits*. New York: Oxford University Press.