

I LOGIC AND RATIONAL BELIEF

1.1 Logic and Reason

IF there is one respect in which humans differ most fundamentally from the other animals, perhaps it is our superior ability to reason about, and understand, the world. The main product of our reasoning, and medium of our understanding, is, of course, also our chief representation of the world—our system of beliefs.

Two dimensions of evaluation come immediately to mind in evaluating a person's beliefs. The first, and most obvious, is accuracy. Beliefs can represent the world more or less accurately, it seems, and the more accurate they are, the better. But philosophers have long been interested in a distinct dimension of evaluation. Some beliefs are more *rational* than others. And though the dimensions of accuracy and rationality may well be linked, they are evidently not the same. A fool may hold a belief irrationally—as a result of a lucky guess, or wishful thinking—yet the belief might happen to be accurate. Conversely, a detective might hold a belief on the basis of careful and exhaustive examination of all the available relevant evidence—in a paradigmatically rational way—and yet the evidence might happen to be misleading, and the belief might turn out to be way off the mark.

The point of evaluating rationality, as well as accuracy, of beliefs surely has to do with our desire to assess the agent qua believer. In doing so, we try to abstract from a certain kind of luck, or accidentalness. The fool is no better a thinker for having guessed correctly. The detective is no poorer a thinker for having

encountered misleading clues. Rational beliefs, it seems, are those arising from good thinking, whether or not that thinking was successful in latching on to the truth.

But what is it that makes thinking “good”? A standard answer to this question is that, at least in part, good thinking is logical thinking. Thus logic has been at the center of philosophical thought about rationality since the time of the ancient Greeks. And the discipline of logic as practiced today incorporates at its center exactly the division between good thinking and accuracy mentioned above. The first lesson in most logic courses—and in many general introductions to philosophy—distinguishes soundness from validity. The latter, which is often thought of as the central concern of the logician, abstracts away from issues of actual truth and falsity to concentrate on studying correct and incorrect relations between claims—relations that are thought to be embodied in good and bad thinking, respectively. Of course, there may turn out to be a deep connection between considerations of truth and considerations which make certain relations between claims “logical.” But the first concern of the logic teacher is typically to separate logical relations from factual ones.

The motivations for studying the logic of rational belief are undoubtedly various. We might seek to improve the thinking of others, or of ourselves, by providing rules that people could self-consciously employ in forming or revising their beliefs. We might seek to diagnose confusion in cases where our thinking naturally leads us to paradoxical results. Or we might simply seek a purely theoretical understanding of good thinking, for its own sake. But regardless of motivation, something like the following idea seems to be presupposed in studying logic: that the correct logic will provide a model for ideally rational belief. It is this idea that I would like to examine and, eventually, defend.

‘Logic,’ of course, is not a term that is used consistently, even within academic philosophy. Texts on logic discuss issues ranging from Gödel’s incompleteness theorem to ways of identifying *ad hominem* arguments. What I have in mind is *formal* logic. Of course,

even the meaning of ‘formal’ is not clear. Texts often refer to, e.g., entailments whose validity depends on logical form or structure; but in explaining what counts as form, the texts typically resort to examples. Thus, the fact that the sentence

(1) Sulfur is yellow

entails the sentence

(2) Sulfur is yellow or sulphur is green

might be cited as flowing from the fact that sentences of the form “P or Q” are always entailed by sentences of the form “P.” This is a paradigmatic principle of formal logic. However, the claim that (1) entails

(3) Sulfur is not red

might well not be considered to flow from any fact about logical form. This is so even though one might claim that sentences of the form “x is not red” are always true when sentences of the form “x is yellow” are. And most would reject the claim that

(4) This rock is made of sulfur

entails

(5) This rock is made of the element with atomic number 16

as a matter of logic, even though many would hold that sentences of the form “x is made of sulfur” can only be true at possible worlds where sentences of the form “x is made of the element with atomic number 16” are true as well.

I do not propose here to settle the questions of whether the second of the above-mentioned entailments is a logical one, or why the third is not. Provisionally, let us take formal logic as concerned with at least the forms or structures created by the standardly accepted logical words such as ‘not,’ ‘or,’ ‘and,’ ‘if...then,’ ‘all,’ and ‘some.’ Could the correct logic of such forms (which is perhaps not the whole of formal logic)

provide a model (undoubtedly a partial model) for ideally rational belief?

Before beginning to answer this question, it will be useful to clarify what sort of rationality is at issue, in several respects.

1.2 Pragmatic vs Epistemic Rationality

A distinction is often made between two senses of rationality, both of which can be applied to beliefs. The distinction is most easily illustrated with a touched-up version of Pascal's Wager. Suppose that, given the evidence available to me, it's unlikely that God exists. However, suppose that the evidence also makes it very likely that if God does exist it will be overwhelmingly in my best interests to toe the theistic line—not only in my actions, but in my beliefs. It could then be rational for me, in the pragmatic sense, to believe in God: given what I want, having that belief could be expected to be most advantageous relative to my ends.¹ But there is also a clear sense in which a belief adopted counter to the evidence would not be a rational one. It is this second, epistemic, sense of rationality that I am concerned with here.

This is not to deny that the two forms of rationality are connected. One might suggest that epistemically rational beliefs are those that would maximize one's expectation of reaching one's particularly epistemic goal or goals (such as believing true things or failing to believe false things). As it stands, the suggestion seems clearly wrong. After all, one can imagine a variant on the Pascal example in which the rewards for the counter-evidential belief were purely epistemic. Consider, for example, a case in which three-fingered aliens offer one vast new insight into physics,

¹ I should note that pragmatic rationality need not be tied to self-interest. It would have to be specified that, *given all my values*, including, e.g., any value I placed on my believing the truth and any value I (perhaps altruistically) placed on other things that could be affected by my beliefs, the expected value of my believing would be higher.

including both new information and corrections of many current misconceptions, on the condition that one believe that Ghengis Khan had three fingers (I'm assuming that the aliens' offer gives one no evidence for this proposition). In this case, it seems clear to me that if one somehow comes to believe that Ghengis Khan was three-fingered, this belief is epistemically irrational—whether or not the aliens can rationally be expected to come through with their part of the bargain. But there are more sophisticated ways of connecting pragmatic and epistemic rationality, some of which will be discussed in later chapters. At this point, I just want to distinguish our epistemic notion from the pragmatic one.

1.3 Diachronic vs Synchronic Rationality

Logic books are often written as though their central topic were inference. Arguments are set out in premise–conclusion form, and it is suggested that the premises represent an agent's present beliefs, and the conclusion a further belief that the agent should, after going through the argument, adopt. This suggests that the sort of rationality being addressed through logic is *diachronic* rationality. Diachronic rationality constrains the way beliefs are changed (or maintained) across time. The idea seems to be something like this: logic's rules of inference tell you which new beliefs you should adopt, on the basis of your current beliefs. Thus Modus Ponens tells you that if you believe P , and you also believe $(P \supset Q)$, then if you don't already believe Q , you should adopt the belief that Q .

It is well-known, of course, that this picture is too simple. After all, if you learn by the above logical argument that your beliefs entail Q , you might well want to revise your current belief that P , or your belief that $(P \supset Q)$, rather than adopt the new belief that Q . This is especially clear if you happen already to believe not- Q . But even if you're currently agnostic on Q , Modus Ponens itself provides no reason for preferring, e.g., becoming a Q -believer to

becoming a P-agnostic.² In each of these cases, logic gives you no guidance at all regarding which option for revising your beliefs is preferable. Thus the “rules of inference” given in logic books cannot be thought of in any straightforward way as rules of diachronic rationality.³

In fact, despite the way logic texts often present the subject, it is far from clear that the basic concern of logic is with change of belief. A valid argument, as most often defined, is one in which it is impossible for the premises to be true while the conclusion is false. In other words, the notion of valid argument flows from a deeper notion, a notion involving the possibility of sentences simultaneously having certain truth-values. Facts about the possible truth-value distributions among the members of a certain set of sentences are not diachronic facts about those sentences. Of course, the facts about possible truth-value distributions have implications for inferences constructed from those sentences—e.g. that certain inferences are truth-preserving. But these implications do not really go beyond the fundamental facts about simultaneous truth-value distributions.

The view of logic as concerned most basically with simultaneous truth-value distributions suggests simpler ways of applying logic to the theory of rational belief. We saw above that logic did not help the rational agent choose between, on the one hand, adopting the

² I should note that, as Mark Moyer pointed out to me, in a typical case where you believe P and $(P \supset Q)$ but are agnostic on Q, you will have evidence for the first two beliefs, and this might well give you good reason for adopting Q over abandoning P. But there are certainly cases in which abandoning P is the rational choice. For example, suppose that the general practice has been for Sophie to wash the dinner dishes. But this morning, I saw Sophie and Sam decide that tonight’s dishwashing duties would be decided by flip of a coin: Sophie will wash iff the coin comes up heads. Nevertheless, not having this memory at the forefront of my mind, I now believe that Sophie will wash the dishes tonight. I also believe (though I don’t bring this belief to mind at the same time as the last one) that Sophie will wash the dishes tonight only if the coin will come up heads. So I believe S and I believe $(S \supset H)$, but I haven’t put two and two together, and I don’t believe H. In this case, when I do notice that my beliefs fail to respect Modus Ponens, it seems that I should become an S-agnostic rather than beginning to believe in H.

³ Harman (1970, 1986) makes this point nicely. In the latter work, he draws the more dramatic conclusion that “there is no clearly significant way in which *logic* is specially relevant to reasoning” (1986, 20).

belief that Q and, on the other hand, dropping the belief that P or the belief that $(P \supset Q)$. But logic did suggest some constraints on our agent's beliefs. In fact, the apparent need for *some revision or other* stemmed from the idea that there was something wrong with certain sets of *simultaneous* beliefs. Most plausibly, it may be claimed that logic precludes the option of rationally believing all of P , $(P \supset Q)$, and not- Q at the same time. And one might also see logic as precluding the option of rationally remaining agnostic on Q while believing both P and $(P \supset Q)$. Generalizing these two suggestions yields the two most prominent proposals for using logic to constrain rational belief: the requirement that a rational agent's beliefs be *logically consistent*, and the requirement that the rational agent's beliefs be *closed under deduction* (i.e. that the agent believe everything logically implied by her beliefs). Both of these proposals are, of course, for synchronic constraints on rationality. Thus I would like to concentrate on the question of whether logic can provide synchronic constraints on ideally rational beliefs.

This may seem misguided, if we want rational beliefs to be those that are in some sense the products of good thinking. "Good thinking" seems to be an activity, and a notion of rationality that tried to capture what good thinking was might seem to be an essentially diachronic notion, involving the evaluation of how we change (or maintain) beliefs through time. Now some have argued that epistemic rationality is fundamentally independent of diachronic considerations.⁴ But even if we put this issue aside, it seems clear that synchronic constraints could play a role in a notion of rationality that was tied to thinking well. For example, "maintain logical consistency" is certainly a rule for how to think, but the constraint it places on ideally rational beliefs concerns how beliefs at a given time relate to other beliefs at that same time. Thus the importance of synchronic constraints is perfectly compatible with conceptions of epistemic rationality according to which it has important dynamic dimensions.

⁴ See e.g. Foley (1993).

1.4 Local vs Global Rationality

One of the first things that is intuitively apparent when one begins to think about rational belief is that some of an agent's beliefs may be more rational than others. For example, although I may be highly responsive to the evidence when I form beliefs about the character, talents, and physical attractiveness of most children I happen to meet, my beliefs about how my own children fare along these dimensions may be subject to considerable non-evidential influences. Clearly, a complete study of rationality must include those factors that affect the rationality of an agent's beliefs differentially. And of course, much work in epistemology has focused on describing what it is about an individual belief that makes it rational, or, more often, what makes it justified, or an instance of knowledge.

This approach encounters a well-known difficulty when it examines questions relating to how structural or logical factors influence epistemic merit. If we are asking, for example, whether my belief that Q is made rational in part by its following logically from my beliefs that P and $(P \supset Q)$, the answer to our question will in part depend on the epistemic status of the latter beliefs. In general, the merits (or demerits) a belief receives in virtue of its logical connections to other beliefs will depend on two entirely separate factors: the nature of the structural relations themselves, and the epistemic credentials of the related beliefs. And evaluating the second component incurs obvious regress problems, since the epistemic credentials of the related beliefs will depend on their structural relations to still other beliefs, whose own epistemic credentials then become relevant; and so on.

There are, of course, various solutions to the regress problem, and this is not the place to examine them. But we may note that focusing on structure leads us to see the rationality of an individual belief as tied into a complex web of logical relations. In short, there's a pressure toward the sort of holism that takes the entirety

of a person's belief system as at least potentially relevant to the epistemic status of any given belief. I don't want to argue here that this degree of holism is correct. But it does seem clear that, insofar as logical relations are important to a belief's epistemic status, they are going to involve large sets of interrelated beliefs.

This is not to say that even a strong form of holism would preclude our allowing differences in epistemic merit to exist among the beliefs of a single agent. But it does make the task of accounting for these differences more complex. And the notions invoked in such an account will naturally be notions of degree—e.g. the degree of a given belief's coherence with other beliefs, or the directness of the connections between the given belief and relevant other beliefs— notions that are unlikely to find precise formal characterization. Thus, insofar as we can capture the factors that are responsible for the differing degrees of rationality enjoyed by various of an agent's individual beliefs, they will likely involve complex and non-formal notions.

This suggests that in studying formal constraints on rationality we might best begin by focusing not on differences in epistemic status among individual beliefs, but rather on the possibility of giving formal constraints on the whole set of an agent's beliefs. And indeed, the initially attractive formal constraints mentioned above—deductive consistency and deductive closure—are global constraints of just this sort. The same holds for probabilistic formal constraints on degrees of belief that have been advanced by many writers.

This global approach to formal constraints on rationality dovetails nicely with seeing such constraints as aspects of an epistemic *ideal*. Concentrating on ideal rationality allows us to put aside, for the moment at least, the difficult questions involved in balancing those non-ideal factors responsible for differences in rationality among the beliefs of a realistic agent.

This approach does not, of course, involve any claim to the effect that all departures from ideal rationality are on a par. For instance, in his coherence-based account of empirical justification, Laurence

BonJour holds that an inconsistency anywhere in an agent's beliefs diminishes, to some extent, the justification of every one of her beliefs (since no belief can cohere perfectly with the entire corpus). BonJour offers no detailed way of assessing how much the rationality of the various beliefs in an agent's corpus is affected by different sorts of inconsistency, but he suggests that different sorts of inconsistency would diminish justification to different degrees. Still, BonJour holds that ideal justification requires that the whole set of the agent's beliefs be logically consistent.⁵

Thus it should be clear that, in concentrating on ideal rationality, we are deliberately putting aside important questions. Real people will presumably always fall below the ideal standards, and any complete account of epistemic rationality must eventually give us insight into the degrees of sub-ideal rationality we find in real people's beliefs. But the fact that a correct description of ideal rationality would not exhaust epistemology does not vitiate the interest of such a description.

Nevertheless, this facile observation—that one cannot do everything at once in epistemology—leaves open some serious worries about both the degree and the nature of idealizations involved in proposals for imposing formal constraints on rational belief. Some of these worries are specific to particular formal models, while others involve the very notion that rationality can be tied to logical standards which are out of reach for real people. These questions will be explored in detail in later chapters. For now, I only want to claim that examining global conditions on ideal rationality is a promising place to start in studying the question of what role formal models can play in understanding rational belief.

To sum up: Our main question is whether formal models have an important role to play in understanding rational belief. This question is complicated by the fact that there are many aspects and varieties of rationality that might be thought to apply to belief, and there can be different approaches to studying a given aspect or

⁵ See e.g. BonJour's "Replies and Clarifications," in Bender (1989, 284).

variety of belief-rationality. My aim has been to focus on that part of the theory of rational belief in which formal models seem most likely to play a useful part. Let us begin, then, by examining formal constraints that apply to the whole set of beliefs an agent has at a given time, and asking whether such constraints can provide a model for ideal epistemic rationality.

2 TWO MODELS OF BELIEF

2.1 Models of Belief and Models of Rationality

WHEN people talk about the world, they typically make unqualified assertions. In ordinary contexts, it is natural to take those assertions as reflecting beliefs of the speaker; if a speaker says “Jocko cheated” (as opposed to “Jocko probably cheated,” or “Jocko must have cheated”), we infer that she bears a fairly simple relation to the claim that Jocko cheated—she believes it. This relation often does not seem to be a matter of degree; either one believes that Jocko cheated, or one doesn’t.¹

Similarly, when people talk explicitly about their beliefs, they often seem to presuppose an all-or-nothing notion. Questions such as “Do you believe that Jocko cheated?” oftentimes seem unproblematically precise. The model of belief that seems implicit in these cases is black-and-white: belief is an attitude that one can either take, or fail to take, with respect to a given claim.

Of course, one also may disbelieve a claim, which is clearly a different thing from failing to believe it, despite the fact that it is natural to express, e.g., disbelief in the claim that Jocko cheated by saying “I don’t believe that Jocko cheated.” But disbelief need not be seen as a third attitude that one can take to a claim. Disbelieving a claim is naturally understood as believing the claim’s negation. Failing to believe either a claim or its negation seems naturally to be expressed by assertions such as “I don’t know whether Jocko

¹ Unqualified assertions may indicate something more than belief, such as claims to knowledge. (One may react to a challenge to one’s unqualified assertion by saying, e.g., “Well, I *believe* that Jocko cheated.”) But even on this stronger reading of what assertions indicate, they seem to indicate a state that includes an all-or-nothing state of belief.

cheated.” So the model of belief that seems implicit in much ordinary thought is naturally taken to be a binary one.²

A binary model of belief also fits in very naturally with philosophical analyses of knowledge. Knowledge has typically been seen as belief-plus-certain-other-things. The belief part has typically been taken as unproblematic—either the agent believes the claim or she doesn’t—and the main task of the theory of knowledge has been taken to be that of providing an adequate specification of what, besides belief, knowledge requires. Even those epistemologists who concentrate on the justification of belief—a topic close to our own—have often seen justification as one of the things a belief needs in order to count as knowledge. Thus mainstream epistemologists of various persuasions have typically employed a binary model of belief.

Nevertheless, the binary model does not provide the only plausible way of conceiving of belief. It is clear, after all, that we have much more confidence in some things we believe than in others. Sometimes our level of confidence in the truth of a given claim decreases gradually—say, as slight bits of counterevidence trickle in. As this occurs, we become less and less likely to assert in an unqualified way (or to say unqualifiedly that we believe) the claim in question. But reflection on such cases fails to reveal any obvious point at which belief suddenly vanishes. At no time does there seem to occur a crisp qualitative shift in our epistemic attitude toward the claim. This suggests that underlying our binary way of talking about belief is an epistemic phenomenon that admits of degrees.

Degrees of belief reveal themselves in numerous ways other than in our introspection of different levels of confidence. Famously, in confronting practical problems in life, whether about what odds to bet at or about whether to carry an umbrella when leaving the

² One might quite reasonably want to avoid equating disbelief in P with belief in P’s negation. In that case, one would naturally see discrete belief as a trinary notion, encompassing three distinct attitudes one might take toward a proposition: belief, disbelief, and withholding judgment. Since nothing relevant to the present discussion turns on the difference between these ways of understanding discrete belief, I will continue to speak of the “binary” conception.

house, our decisions and actions seem to be explained by degrees of belief. Rational explanations of an agent's actions typically make reference to the agent's beliefs and desires. The desire-components of such explanations obviously depend not only on the contents of the agent's desires, but on their strengths. And similarly, the belief-components of such explanations depend on the agent's degrees of confidence that the various possible choices open to her will lead to outcomes she cares about. The common-sense psychological principle that underlies these explanations seems to be a rough approximation of expected utility maximization: in the textbook umbrella case, for example, the greater an agent's confidence that leaving the umbrella at home will result in her getting wet, and the more strongly she disvalues getting wet, the less likely she will be to leave the umbrella at home. Thus, a sizable minority of epistemologists have approached the rationality of belief from a perspective closely intertwined with decision theory, a perspective in which degrees of belief are taken as fundamental.

Both the binary and the graded conceptions of belief enjoy, I think, at least a strong *prima facie* plausibility. And each conception figures in apparently important philosophical thought about rationality. Thus, although it could turn out in the end that one (or both) of these conceptions failed to pick out any epistemically important phenomenon, we should not dismiss either one at the outset as a potential home for formal rationality requirements. Still, this leaves open a number of possible approaches to the objects of epistemic rationality. One might see binary belief as reducing to graded belief, or graded belief to binary belief. In such a picture, there would be at bottom only one fundamental object of rational appraisal. Alternatively, one might see two independent (though undoubtedly related) epistemic phenomena. In this case, perhaps each would be answerable to its own distinctive set of rational demands.

Getting clear on this issue is important to our purposes, because the two conceptions of belief seem to invite quite different kinds of formal models. The traditional binary conception of belief meshes

naturally with straightforward applications of deductive logic. On the binary conception, there is a set of claims that a given agent believes. The basic idea is, roughly, that membership in this set of claims ought (ideally) to be conditioned by the logical properties of, and relationships among, those claims. As we've seen, deductive consistency and deductive closure are prominent candidates for constraints on an ideally rational agent's set of binary beliefs.

By contrast, the graded conception of belief requires quite a different treatment. On this conception, there is not one distinctive set of claims the agent "believes"; instead, the agent takes a whole range of attitudes toward claims. At one end of the spectrum are those claims the agent is absolutely certain are true, at the other end are claims the agent is absolutely certain are false, and in between are ranged the vast majority of ordinary claims, in whose truth the agent has intermediate degrees of confidence. The standard formal models for ideally rational degrees of belief involve using the probability calculus. Degrees of belief are taken to be measurable on a scale from 1 (certainty that the claim is true) to 0 (certainty that the claim is false). An ideally rational agent's degrees of belief must then obey the laws of probability; to use the common terminology, they must be probabilistically coherent.

The probability calculus is often referred to as a logic for degrees of belief. It might be more illuminating to see it as a way of applying standard logic to beliefs, when beliefs are seen as graded. The constraints that probabilistic coherence puts on degrees of belief flow directly from the standard logical properties of the believed claims. Consider, for example, the fact that probabilistic coherence requires one to believe $(P \vee Q)$ at least as strongly as one believes P . This flows directly from the fact that $(P \vee Q)$ is logically entailed by P . In fact, we can plainly see connections between the natural ways logic has been taken to constrain belief on the binary and graded conceptions. The dictate of logical closure for binary beliefs requires that

an ideally rational agent does not believe P while failing to believe $(P \vee Q)$.

Probabilistic coherence of graded belief requires that

an ideally rational agent does not believe P to a given degree while failing to believe $(P \vee Q)$ to at least as great a degree.

Similarly, logical consistency of binary belief requires that

an ideally rational agent does not believe both P and $\sim (P \vee Q)$;

in other words, if she believes one of the two sentences, she does not believe the other. Probabilistic coherence of graded belief requires that

an ideally rational agent's degrees of belief in P and $\sim(P \vee Q)$ do not sum to more than 1;

in other words, the more strongly she believes one of the two sentences, the less strongly she may believe the other.

The idea that the probability calculus functions less as a new logic for graded belief than as a way of applying our old logic to graded belief may be supported by looking at the basic axioms of the probability calculus. Put informally, they are as follows (where $\text{pr}(P)$ stands for the probability of P):

- (1) For every P , $\text{pr}(P) \geq 0$.
- (2) If P is a tautology, then $\text{pr}(P) = 1$.
- (3) If P and Q are mutually exclusive, then $\text{pr}(P \vee Q) = \text{pr}(P) + \text{pr}(Q)$.

The above formulation is quite typical in using the notions of tautology and mutual exclusivity. These notions are, of course, the standard logical ones. Presentations of the second axiom sometimes use "necessary" rather than "a tautology," but insofar as necessity and logical truth come apart, it is the latter that must be intended. No one thinks, presumably, that the axioms of probability should be applied to rational belief in a way that requires "Cicero is Tully" to have probability 1.

This observation suggests that the import of the standard axioms is parasitic on a pre-understood system of deductive logic. On any system of logic, $(P \ \& \ Q)$ will entail P , and this will be reflected directly in restrictions on probabilistically coherent degrees of beliefs that one may have in these propositions. But the boundaries of logic are not entirely obvious. If it is a matter of logic that $\Box P$ entails P , then $(\Box P \supset P)$ will be a tautology, and $\Box P$ and $\sim P$ will be mutually exclusive, and this will in part determine which degrees of belief involving these sentences can be probabilistically coherent. Similarly, when we decide whether, as a matter of logic, $\Box P$ entails $\Box \Box P$, or “ x is yellow” entails “ x is not red,” or “ x is made of sulfur” entails “ x is made of the element with atomic number 16,” we will thereby determine the contours of probabilistic coherence. That is why the axioms of probability are better seen not as a distinct logic for graded beliefs. The probability calculus is most naturally seen as just giving us a way of seeing how rational graded beliefs might be subject to formal constraints derived directly from the standard logical structures of the relevant propositions.

Now it is true that there are ways of axiomatizing the probability calculus that do not separate the probabilistic axioms from those of deductive logic. For example, Karl Popper (1959) gives an axiomatization for conditional probability that incorporates standard propositional logic (he intends it as a generalization of deductive propositional logic). Hartry Field (1977) extends Popper’s technique to give an axiomatization that incorporates predicate logic (Field intends not to generalize deductive logic, but rather to provide a truth-independent semantics which reflects conceptual roles rather than referential relations).³ We should be careful, then, about what we conclude from examining standard formulations of probability theory (or the formulations used by the theory’s developers): even if the standard axiomatizations are intuitively natural, that does not prove that the probability calculus is, at the most fundamental level, parasitic on a conceptually prior system of deductive logic.

³ The relevance of this point was brought to my attention by a referee. See also Hawthorne (1998) for further development and related references.

However, the point remains that probability theory is in no way independent of the ordinary logical relations familiar from deductive logic—relations that derive from important structural patterns involving ‘and,’ ‘not,’ ‘all,’ etc. The constraints that any version of probability theory places on degrees of belief flow from exactly these patterns. And the standard way of axiomatizing probability shows that, for any of the familiar notions of deductive consistency, there will be a probabilistic way of taking account of that logic’s structural basis.

For both models of belief, then, the prominent proposals for imposing formal constraints on ideal rationality are rooted in logic. But the logic-based constraints take quite different forms for the different models of belief. Moreover, it turns out that the arguments both for and against the imposition of the formal constraints are quite different for binary and graded belief. Thus our examination of the plausibility of formal constraints on rational belief will clearly be shaped by our choice of how to see rational belief itself.

2.2 Unification Accounts

We saw above that both conceptions of belief enjoy enough plausibility to be worth exploring, and thus that we should not reject either out of hand. But even putting aside the eliminationist option of rejecting one of the conceptions as not picking out any real phenomenon, one might favor what might be called a unification approach. One might hold that one sort of belief was really only a special case or species of the other. If such a view were correct, it clearly could help determine our approach to formal rationality.

Perhaps the less attractive unificationist option is to take graded beliefs as nothing over and above certain binary beliefs. Let us consider an example in which a graded-belief description would say that an agent had a moderate degree of belief—say, 0.4—in the

proposition that Jocko cheated on Friday's test. Should we see this graded belief as really consisting merely in the agent's having some particular binary belief? If so, we should presumably turn our attention straightforwardly to deductive constraints.

The problem with this proposal stems from the difficulty of finding an appropriate content for the relevant binary belief. A first try might be that the probability of Jocko's having cheated on Friday's test is 0.4. But what does "probability" mean here? The term is notoriously subject to widely divergent interpretations. Some of these interpretations—those of the "subjectivist" variety—define probability explicitly in terms of graded belief. Clearly, if graded beliefs are merely binary beliefs about probabilities, the probabilities involved must not be understood this way.

On the other hand, if we understand probabilities in some more objective way, we risk attributing to the agent a belief about matters too far removed from the apparent subject matter of her belief. For example, if probabilities are given a frequency interpretation, we will interpret our agent as believing something like: Within a certain specific reference class (cases where people had a chance to cheat on a test? cases where people like Jocko had a chance to cheat on a test? cases where Jocko himself had a chance to cheat on a test on a Friday? ...), cheating took place in 4/10 of the cases. Yet it is hard to believe that any thought about reference classes need even implicitly be present in the mind of an agent to whom we would attribute a 0.4 degree of belief in Jocko's having cheated. If probability is given a propensity interpretation, things are no better. Since the belief in question is about a past event, we cannot say that the agent believes that some current setup is disposed to a certain degree to end up with Jocko cheating on the test in question. And it seems quite implausible to analyze our agent's belief as really being about the way Jocko was disposed to behave at a certain point just prior to the test.

One could object to this argument that precise degrees of belief are almost never correctly attributable, and that my example therefore should not have specified a degree as specific as 0.4 in the first

place. The agent, it might be held, really only harbored a (binary) belief that Jocko's cheating was quite possible, but not highly probable. But while there may be some point behind the charge that the attribution of precisely a 0.4 degree of belief in this case is unrealistic, softening the focus here to talk about more vague probability-beliefs does not address the present worry. The worry, after all, was that when people have intermediate degrees of belief in propositions, they need not have any beliefs at all about, e.g., frequencies within reference classes, or propensities.

Of course, these examples are based on quick and crude caricatures of prominent objective interpretations of probability, and still other objective accounts of probability do exist. But for our purposes, these examples serve well enough to show how unnatural it is to identify an agent's having a certain degree of confidence in a particular proposition with that agent's having an all-or-nothing belief about some non-belief-related proposition about objective probabilities.

Moreover, it is clear that, in general, people's attitudes do come in degrees of strength. Presumably, no one would doubt the existence of degrees of strength with respect to people's hopes, or fears, or attractions, or aversions. Yet on the unification view about belief that we have been considering, strength of confidence would have no reality independent of (binary) beliefs about objective probabilities. I see little reason to accept such a view. So although this sort of unification would simplify matters by turning our attention to deductive, as opposed to probabilistic, constraints on rational belief, it seems unlikely that trying to simplify matters in this way would be successful.

A more promising sort of unification would work in the opposite way. We might see binary belief as a special case or species of graded belief: one would believe something in the binary sense if she believed it (in the graded sense) with a strength that met a certain threshold. Two variants of this proposal have in fact been advanced. According to one, binary belief is identified with graded belief of the highest degree (1); on this account, to believe P is to be

certain that P. According to the other account, the threshold is lower (and may not be precisely specified); on this account, to believe P is to be sufficiently confident, but not necessarily certain, that P. Let us consider these accounts in turn.

The certainty proposal is, I think, less plausible. If the binary conception of belief derives its plausibility from our habit of making unqualified assertions, and from our ordinary ways of thinking and talking about belief, then the plausible notion of binary belief is of an attitude that falls far short of absolute certainty. We often assert, or say that we believe, all kinds of things of which we are not absolutely certain. This is particularly clear if the plausibility of the graded conception of belief is rooted in part in how belief informs practical decision. Insofar as degree of belief is correlated with practical decision-making, the highest degree of belief in P is correlated with making decisions that completely dismiss even the tiniest chance of P's falsity. For example, having degree of belief ι in Jocko's having cheated would correlate with being willing literally to bet one's life on Jocko's having cheated, even for a trivial reward. Surely this level of certainty is not expressed by ordinary unqualified assertions; nor is it what we usually want to indicate about ourselves when we say, e.g., "I believe that Jocko cheated," or what we want to indicate about others when we say, e.g., "Yolanda believes that Jocko cheated."

Now one might resist taking too strictly our everyday tendencies to attribute belief in cases such as Jocko's cheating, and still insist that there is an important class of ordinary propositions about the external world which we rationally accord probability 1. Isaac Levi (1991) has argued that we do, and should, have this sort of "full belief" even in propositions that we come to believe by methods which, we recognize, are not absolutely reliable. When we accept such propositions as evidence, we "add [them] to the body of settled assumptions," which are "taken for granted as settled and beyond reasonable doubt" (1991, 1). According to Levi, these propositions then function as our standard for "serious" (as opposed to merely logical) possibility.

However, it does not seem to me that we are actually fully certain even of the things we typically take for granted or treat as evidence. It is, of course, true that there are many propositions which, in some rough sense, we regard as settled in our practical and theoretical deliberations. For example, scientists studying the effects of a new drug on rats may accept as evidence a proposition such as

The rats treated with drug D died, while the rats in the control group lived.

In evaluating hypotheses about the drug, the researchers will consider various explanations for this evidence—that drug D caused the deaths of the treated rats; that the batch of saline solution in which drug D was dissolved contained a contaminant that caused the deaths of the treated rats; that it was just a coincidence; etc. But they will not consider the possibility that the evidence proposition is actually false. In an ordinary sense, this possibility will not be taken as “serious.”

Does this mean that the researchers are absolutely certain of the evidential proposition? I don’t think so. We would not, for example, expect one of them to be willing to bet the lives of his children against a cup of coffee on the proposition’s truth. And we would not think that it would be reasonable for him to do this. Why? Because there is some incredibly small chance that, e.g., the lab technician switched the rats around to make the experiment “come out right.” What would explain the researcher’s reluctance to take the bet (or our reluctance to call the bet reasonable) is precisely the fact that the researcher is not completely certain of the evidential proposition.

But let us put this sort of doubt aside, and consider the consequences of accepting a unification account on which binary belief was identified with graded belief of probability 1. It remains true that the graded conception of belief has within it the notion of “full belief,” or belief with degree 1. And one might argue for a kind of unification (perhaps one that deviated from some aspects of our intuitive conceptions) by identifying binary belief with full belief.

If we were to accept this sort of unification, what impact would it have on the question of formal constraints on rational belief?

Clearly, the fundamental approach to rational constraints would be the one appropriate to graded belief—presumably, a probabilistic one. And adopting such an approach would actually automatically impose constraints on binary belief—in fact, constraints that would at least come close to the traditional deductive constraints of consistency and closure.⁴ But the status of the (approximation to the) traditional deductive constraints on this picture would be derivative. Insofar as the certainty proposal is plausible, then, it argues for taking a probabilistic approach to formally constraining rational belief.

Perhaps, however, it is more plausible to unify the two conceptions of belief by setting the binary belief threshold at some level below that of certainty. One needn't hold that our ordinary notion picks out some precise cutoff value (“if it's believed to at least degree 0.9, it is Believed”); one might hold instead that the border of binary belief is a vague one. Still, one might develop a model of rational belief that incorporated a precise (if somewhat arbitrary) cutoff point, in order to study the formal constraints that might apply on any such precisification.

This sort of unification comes closer than does the certainty proposal to fitting with our ordinary practices of unqualified assertion and belief-attribution. By and large, it seems, we do make assertions and attribute (binary) beliefs in cases where degrees of

⁴ The constraints imposed on full beliefs by the probability calculus coincide with those imposed on binary beliefs by traditional consistency and closure conditions in many ways. For example, one cannot fully believe a contradiction; one must fully believe tautologies; one cannot have less than full belief in $(P \vee Q)$ while having full belief in P ; and one cannot have full belief in all of P , $(P \supset Q)$, and $\sim Q$. The divergences can occur in certain contexts involving infinite sets of beliefs. For example, if one is certain that something is located at a point somewhere in a given area, but thinks that all the infinite number of points in the area are equally likely, it turns out that the probability assigned to the thing being at any one point must be 0, and hence the probability of it not being at that point must be 1. Thus one must have full belief that the thing is not at p , for each point p in the area—even though one also has full belief that the thing is at one of these points. In this sort of case, then, one has an inconsistent (though not finitely inconsistent) set of beliefs. See Maher (1993, ch. 6.2) for detailed discussion of this matter.

belief are fairly high. Thus, of all the unification proposals considered so far, this one may be the most likely to be correct.⁵

On this sub-certainty threshold account, it is not true that imposing probabilistic constraints on graded belief automatically imposes deductive-style constraints on binary belief. There's no reason to think, for example, that the set of things a rational agent believes to at least degree 0.9 should be consistent with one another. In fact, quite the reverse is true, for any sub-certainty threshold, as is made clear by lottery examples. (Consider a rational agent who has excellent evidence, and is thus very highly confident (> 0.999), that a particular 1,000-ticket lottery is fair, and that one of its tickets will win. For each ticket, his confidence that it won't win is 0.999. Thus he is rationally confident, to an extremely high degree, of each member of an inconsistent set of propositions.) Henry Kyburg famously used this point in arguing against taking deductive consistency to be a requirement on binary belief.⁶ Others have used it in the opposite way, arguing that since deductive consistency is a rational constraint on binary belief, belief in a proposition cannot simply be a matter of having sufficient confidence in it.⁷

I don't want to take a stand here on whether our ordinary binary conception of belief is best understood as referring to a certain level of confidence. Although our assertion and attribution practices may fit better with this account than with the certainty account, the fit is not perfect, especially in lottery cases.⁸ Still, one might well maintain that our talk of binary belief is most plausibly construed as referring to a high level of graded belief, and then work to explain away tensions with our assertion and attribution practice (e.g. by invoking principles of conversational implicature). How would such an approach affect the question of formal epistemic constraints?

⁵ Foley (1993, ch. 4) provides a clear and detailed defense of this sort of view.

⁶ See his "Conjunctivitis," in Swain (1970).

⁷ For recent examples of this argument, see Maher (1993, ch. 6) and Kaplan (1996, ch. 3).

⁸ The status of our attitudes toward lottery tickets (and related matters) will be discussed in more detail in later chapters.

As noted already, the classical constraint of deductive consistency for binary beliefs would have to be given up. The same would then hold for deductive closure: in standard lottery cases, for example, “no ticket will win” follows deductively from propositions each of which meets the confidence threshold for belief, but it does not come close to meeting that threshold itself. As Kyburg points out, binary belief on such an account could still obey vastly weakened versions of these constraints. Beliefs could obey the “Weak Consistency Principle” requiring that no one belief was a self-contradiction. And they could respect a weak version of deductive closure, the “Weak Deduction Principle,” requiring that anything entailed by a single belief was also believed.

Nevertheless, for our purposes, the important point is that these weak principles are simply automatic consequences of imposing probabilistic coherence on the agent’s graded beliefs. Weak Consistency would follow from probabilistic coherence because contradictions have probability 0, and thus would fall below the threshold. Weak Deduction would follow because any logical consequence of a sentence must have at least as high a probability, so if P meets the threshold and P entails Q , Q must meet the threshold as well.

In fact, Kyburg points out that somewhat stronger consistency principles can be imposed, depending on the threshold chosen. If the threshold is over 0.5, “Pairwise Consistency” follows: no pair of inconsistent propositions may be believed (though an inconsistent triad is not ruled out). And in general, as the threshold for belief becomes higher, increasingly larger sets of jointly inconsistent beliefs will be prohibited. Of course, even at a very high threshold (e.g. 0.99), the system will allow large sets (e.g. 101) of jointly inconsistent beliefs.⁹

⁹ Think of an agent who is extremely confident that a certain 100-ticket lottery is fair; the inconsistent set of beliefs will be 100 particular beliefs of the form “ticket n won’t win,” along with the general belief that one of the tickets will win. See Hawthorne and Bovens (1999) for an interesting and detailed exploration of the sorts of consistency constraints that may be imposed in lottery and related cases, given a threshold model of binary belief.

Does this show that the threshold view makes a place for significant deductive constraints on rational belief? It seems to me that it does not. For one thing, it is not clear why “n-wise consistency” principles should be intuitively attractive, from the point of view of describing ideal rationality. Of course, there is intuitive reason to impose the probabilistic constraints on graded belief upon which the limited-consistency principles supervene. But considered apart from the probabilistic constraints, there’s nothing attractive about principles that one can believe inconsistent sets of beliefs only so long as they contain at least n , or at least $n+1$, members.

Moreover, when one moves to consider closure principles, the threshold model does not support similar limited versions of closure. As we’ve seen, one of the motivations for taking deductive constraints seriously is to account for intuitions such as the following:

If an ideally rational agent believes both P and $(P \supset Q)$, she believes Q .

Suppose we tried to advance a limited closure principle as follows: if Q is entailed by any *pair* of an ideally rational agent’s beliefs, then the agent believes Q . This would seem to answer to the intuition above. But it would also amount to imposing an unlimited closure requirement. For any two beliefs will entail their conjunction; and, once that is admitted as a belief, it may in turn be conjoined with a third belief, etc., until the agent is required to believe any proposition that is entailed by any finite number of her beliefs. This is, of course, incompatible with the threshold account of rational binary belief, as the lottery cases demonstrate.¹⁰

Thus it seems that, insofar as sub-certainty threshold accounts of binary belief are plausible, we should look not to deductive constraints, but to probabilistic constraints, if we are to find plausible formal conditions on rational belief. We’ve seen above that a similar

¹⁰ Indeed, the burden of Kyburg’s (1970) “Conjunctivitis” is to cast doubt on the Conjunction rule for rational belief—that if an agent rationally believes P and rationally believes Q she must also believe $(P \& Q)$.

lesson holds for certainty accounts of binary belief. We've also seen that it is not plausible to unify belief by identifying graded beliefs with particular binary beliefs. Summing up, then, it seems that, while no unified account of belief is fully compelling, to the extent that graded and binary belief could be unified, the formal constraints that characterize ideally rational belief would likely be probabilistic.

Still, given that even the threshold account considered above is intuitively problematic, it is worth seeing whether a view of binary belief that made it more independent of graded belief could provide a home for deductive logical constraints. Such a view would, of course, divorce the two kinds of belief in a fundamental way. But several writers have advocated just this sort of divorce.

2.3 Bifurcation Accounts

Bifurcation accounts hold that binary beliefs are different in kind from graded beliefs—that neither is a mere species or special case of the other. Such accounts may be urged for various reasons. For one thing, bifurcation may allow for a better fit with some aspects of our ordinary assertion and attribution practices. In lottery cases in particular, we are reluctant to assert unqualifiedly “This ticket will not win,” even when the lottery is large. Those who would tie binary belief closely with unqualified assertion may take this as important evidence against identifying binary belief with high confidence.¹¹ And there are other cases—in particular, those of apparently rational scientists discussing fairly comprehensive

¹¹ Maher (1993, 134) and Kaplan (1996, 127) explicitly support their bifurcation accounts in this way. Others, however, see assertability as tied to knowledge rather than belief (see Unger 1975, ch. 6; Williamson 1996; and DeRose 1996). DeRose, for example, would attribute belief in lottery cases such as the one described, holding that unqualified assertions would be improper because “this ticket won’t win” would violate a counterfactual tracking-style requirement for knowledge (i.e. you would have the same belief even if you were holding the winning ticket).

theories—when unqualified assertions seem to be made about claims in whose complete truth no one should have very high confidence, given the history of science.¹²

Moreover, it must be acknowledged that even ordinary belief-attributions seem strained in lottery-type situations. Suppose that we know that Yolanda holds a ticket in a lottery she knows to be large, and that she has no special information about her ticket. Suppose we also know Yolanda to be highly rational. We would not hesitate to attribute to Yolanda a high degree of belief in her ticket not winning. But we might hesitate to say, flatly, “Yolanda believes that her ticket won’t win.” And if we asked Yolanda herself “Do you believe your ticket’s a loser?” it would seem at least somewhat unnatural for her simply to reply “Yes.”¹³

If unqualified assertion is taken as a mark of belief, then our ordinary assertion practices also seem to fit uneasily with threshold accounts in a way that is independent of lottery-type cases. Often, our willingness to make unqualified assertions seems to depend on aspects of the context quite independent of the likelihood of the relevant proposition’s truth. Suppose, for example, that ten minutes ago I chatted in my driveway with the neighbors who live on either side of my house, after which I saw them disappear into their respective houses. I know that neither had plans to leave soon, but I haven’t been watching their driveways. Someone knocks on my door by mistake, wanting to speak to my left-hand neighbor about an upcoming concert. I might well say to the person, “Jocko’s at home next door.” On the other hand, when a doctor knocks on my door by mistake, wanting to consult my right-hand neighbor on an emergency life-and-death decision about her relative, I would not say “Yolanda’s at home next door.” I might say that she’s probably at home, or even almost certain to be at home, but I wouldn’t just say unqualifiedly that she was at home. Some have

¹² Maher (1993) argues along these lines; his views on theory acceptance will be discussed in Chapter 4.

¹³ On the other side, though, as DeRose (in correspondence) points out, it would also be unnatural—maybe even more so—for her simply to reply “No.”

used this sort of case to suggest that belief is sensitive to what is at stake in a given matter, and not just to the agent's degree of confidence that the proposition is true.¹⁴

From our perspective, however, the most interesting argument advanced in support of bifurcation accounts is not about fit with ordinary assertion and attribution practices. It is a more theoretical one, which applies directly only to rational (or reasonable, or warranted, or justified) binary beliefs. If the standard deductive consistency and closure constraints apply to rational binary belief, then it cannot be rational to believe that a given large lottery will have a winning ticket, while simultaneously believing of each ticket that it will not win. Now no one seems to want to deny that it can be rational to believe that a big fair lottery will have a winning ticket. But various philosophers have devised conditions on justification, warrant, acceptability, etc., that are expressly intended to preclude rationally believing of any particular ticket that it will lose, no matter how high the odds. If we reject the requirement that rational belief be absolutely certain, it is argued, then only a bifurcated account can possibly allow for binary beliefs to be made subject to rational constraints of deductive consistency and closure. Thus bifurcation views are endorsed precisely because they allow for rational binary beliefs to be governed by logic.¹⁵

Now since the deductive constraints apply only to *rational* beliefs, it might be doubted that their application could be used to argue convincingly for a conclusion about the metaphysics of binary belief in general. And some epistemologists who have defended deductive constraints in the face of lottery examples do not seem to have had metaphysical conclusions explicitly in mind. Bonjour, for example, holds that in lottery cases one does not have a *fully justified belief* that one's ticket will lose. He points out that a belief's degree of justification cannot then be correlated with the

¹⁴ See e.g. Nozick (1993, p. 96 ff.).

¹⁵ For examples of arguments against sub-certainty threshold views of rational belief, see Kaplan (1996, 93 ff.), Maher (1993, 134), Pollock (1983), Lehrer (1974, 190-2), and Bonjour (1985, 54-5).

probability of the belief's truth. But he does not explicitly address the question of whether binary belief itself—the sort of belief with which he is concerned—is an attitude that goes beyond having a certain degree of confidence in the relevant proposition.¹⁶

It is worth seeing, then, whether a unificationist about the metaphysics of belief—say, a sub-certainty threshold theorist—could accommodate the deductive constraints on rational belief. He would have to admit that, when an agent's degrees of belief in the members of the inconsistent set of lottery propositions are each over the threshold, the agent does indeed harbor inconsistent (binary) beliefs. However, he would hold that the beliefs in question were not fully rational (or completely justified, or warranted).

This line seems unpromising to me. Our unificationist must acknowledge that the agent contemplating the large lottery *should* have a high degree of belief in, e.g., the proposition that ticket no. 17 won't win. But if her having a high degree of belief in this proposition is fully rational, and if having the binary belief is nothing over and above having a high degree of belief, then it is surely something of a strain to suggest that the binary belief that ticket no. 17 won't win is not rational in this case. It is, after all, one and the same attitude toward one and the same proposition—that is the essence of the unification approach.

The threshold theorist might try to differentiate between different types of rationality: the agent's attitude might be claimed to be degree-rational but not binary-rational. Surely there is nothing wrong with acknowledging different dimensions of rationality, and admitting cases where they give different verdicts about the same object. For example, one might reasonably think that having a

¹⁶ In Bonjour's description of the belief component of knowledge, there is no obvious mention of any factor going beyond degree of confidence: "I must *confidently believe* ..., must accept the proposition in question without serious doubts or reservations. Subjective certainty is probably too strong a requirement, but the cognitive attitude in question must be considerably more than a casual opinion; I must be thoroughly convinced...." (1985, 4). I should note that this description is part of an account for which he claims only approximate correctness; nevertheless, the reservation he expresses about the belief component is unrelated to the present issue.

certain religious belief, or a belief in the fidelity of one's friend, was pragmatically rational, but that having exactly the same attitude toward exactly the same proposition was epistemically irrational.

Nevertheless, I think that this sort of move will not work in the present case. For in calling an agent's attitude toward a certain proposition irrational one is endorsing a perspective from which the agent's attitude toward that proposition is undesirable. In the present case, since binary-rationality is an epistemic notion, the perspective will have to be an epistemic one. But it is clear that there is nothing at all to be said, from any epistemic perspective, against our agent's high degree of confidence in the proposition that ticket no. 17 will lose. There is no epistemic perspective from which her having a lower degree of confidence would be at all preferable. Thus it turns out that a unifying view cannot accommodate deductive constraints on binary belief by distinguishing degree-rationality from binary-rationality: doing so would deprive binary-rationality of all normative force.

It seems, then, that the plausibility of imposing deductive constraints on rational binary beliefs does have implications for the metaphysics of binary belief in general. Unless we hold binary belief equivalent to certainty, the imposition of the deductive rational constraints requires that binary belief be divorced from graded belief in a fundamental way. Believing a proposition must involve taking some attitude toward it that is wholly distinct from one's confidence that the proposition is true.

Of course, the power of any argument that sought to support a bifurcated metaphysics of belief in this way would depend directly on showing independently that it was plausible to impose the deductive constraints in the first place. Whether this can be done is a question that will be examined closely in the following two chapters. At this point, we can say that a bifurcated metaphysics of belief may find some support in our ordinary assertion and attribution practices, and is a prerequisite to the imposition of the standard deductive constraints on rational belief.

The questions of how, and whether, rational belief is constrained by logic are intimately connected with the question of what belief is. On either a graded or a binary conception, logical relations among propositions can be used to constrain rational belief. But the two conceptions invite quite different ways of doing so: the binary conception invites the imposition of deductive closure and consistency, while the graded conception invites the imposition of probabilistic coherence.

Both conceptions of belief have at least *prima facie* claims to describing important features of our epistemic lives. But the relation between the two kinds of belief is not obvious. Unifying the two conceptions by seeing one kind of belief as a special case or species of the other seems plausible only in one direction (assimilating binary to graded belief). This would leave probabilistic coherence as the fundamental formal constraint on rational belief. In fact, the more plausible route to unification, the sub-certainty threshold approach, is incompatible with taking full-blooded deductive constraints as normative requirements on rational belief. It seems, then, that imposing the deductive constraints requires adopting a fundamentally bifurcated view of belief; the next two chapters will explore this possibility. Probabilistic constraints, on the other hand, may find a home on either a unified or a bifurcated metaphysics of belief; the plausibility of probabilistic constraints will be explored in subsequent chapters.