Chapter 1

Beliefs and Degrees of Belief

Like most epistemological theories, Bayesian Epistemology concerns propositional attitudes. A **propositional attitude** is an attitude an agent adopts towards a proposition, or towards a set of propositions. While much philosophical ink has been spilled over the nature of propositions, we will assume only that a **proposition** is an abstract entity expressible by a declarative sentence and capable of having a truth-value. For example, the sentence "Nuclear fusion is a viable energy source" expresses a proposition. If I believe fusion is viable, that belief is a propositional attitude.

Humans adopt a variety of attitudes towards propositions. I might *hope* that fusion is a viable energy source, *desire* that fusion be viable, *wonder whether* fusion is viable, *fear* that fusion is viable, or *intend to make it the case* that fusion is a viable energy source. While some propositional attitudes involve plans to change the world, others attempt to represent what the world is currently like.

Belief—a propositional attitude central to epistemology—is a representational attitude. (Knowledge, another central representational attitude in epistemology, will not be a major focus of this book.[1]) Belief is in some sense a *purely* representational attitude: when we attribute a belief to an agent, we are simply trying to describe how she takes the world to be, without attributing to her any particular emotional affect towards a proposition, level of justification in that proposition, etc. Yet belief is not the only purely representational attitude; an agent might be *certain* that a proposition is true, or *disbelieve* a particular proposition. Philosophers often discuss the class of **doxastic attitudes** ("belief-like" attitudes) into which belief, disbelief, and certainty fall. Bayesian Epistemology focuses on a doxastic attitude known as degree of belief, degree of confidence, or **credence**.

In recent decades credences have become more prominent in epistemology, as well as in other areas of philosophy (not to mention psychology, economics, and other nearby disciplines). This chapter tries to indicate why that has occurred. I'll begin by contrasting degree of belief talk with other doxastic attitude attributions—especially attributions of so-called "binary" beliefs that have historically been common in epistemology. I'll then examine what working with degrees of belief adds to our account of an agent's doxastic life. Finally I'll introduce a basic characterization of Bayesian Epistemology, and outline how we will explore that view in chapters to come.

1.1 Binary beliefs

1.1.1 Classificatory, comparative, quantitative

In his (1950), Rudolf Carnap helpfully distinguishes classificatory, comparative, and quantitative concepts:

> *Classificatory concepts* are those which serve for the classification of things or cases into two or a few [kinds].... *Quantitative concepts*... are those which serve for characterizing things or events or certain of their features by the ascription of numerical values.... *Comparative concepts*... stand between the two other kinds.... [They] serve for the formulation of the result of a comparison in the form of a more-less-statement without the use of numerical values. (p. 9)

In Carnap's famous example, describing the air in a room as *warm* or *cold* employs classificatory concepts. Characterizing one room as *warmer* than another uses a comparative concept. The *temperature* scale describes the heat of a room with a quantitative concept.

Both our everyday talk about doxastic attitudes and our philosophical theorizing about them use classificatory, comparative, and quantitative concepts. Classificatory terms include *belief*, *disbelief*, *suspension of judgment*, and *certainty*. These doxastic attitudes are adopted towards one proposition at a time; an agent either has the attitude towards a proposition or she doesn't. (So these classificatory attitudes are sometimes referred to as "binary". I'll tend to alternate between the "classificatory" and "binary" terminology in what follows.) A comparative attitude, on the other hand, is adopted towards an ordered pair of propositions. For example, I am *more confident* that fission is a viable energy source than I am that fusion is. A quantitative attitude assigns a numerical value to a single proposition; my physicist friend says she has a 90% *degree of confidence* that fusion is viable.

Until the last few decades,much of epistemology focused around classificatory concepts. (Think of debates about the justification of *belief*, or about necessary and sufficient conditions for *knowledge*.) This was not an exclusive focus, but more a matter of emphasis. So-called "traditional" or "mainstream" epistemologists certainly employed comparative and quantitative terms.[2] Moreover, their classificatory attitude ascriptions could be subtly shaded by various property modifiers: a belief, for example, might be *reluctant*, *intransigent*, or *deeply-held*. Nevertheless, Bayesian epistemologists shifted more of their emphasis to quantitative attitudes like credences.

This chapter tries to explain that shift: Why might an epistemologist emphasize credences over other doxastic attitudes? To help with that explanation I'll introduce a character who I doubt has ever existed in real life: the Simple Binarist. A Simple Binarist insists on describing agents' doxastic propositional attitudes exclusively in terms of belief, disbelief, and suspension of judgment. The Simple Binarist eschews all other doxastic attitude attributions, and even refuses to add shading property modifiers like the ones above. I introduce the Simple Binarist not as a plausible rival to the Bayesian, but instead as an illustrative contrast. By highlighting doxastic phenomena the Simple Binarist has trouble accounting for, I will illustrate the importance of quantitative attitude attributions.

Nowadays most everyone uses a mix of classificatory, comparative, and quantitative doxastic concepts to analyze agents' doxastic lives. I hope to demonstrate the significance of quantitative concepts within that mix by imagining what would happen if our epistemology lacked them entirely. And I will suggest that epistemologists' growing understanding of the advantages of degree-valued doxastic concepts helps explain the preponderance of quantitative attitude ascriptions in epistemology today.

1.1.2 Shortcomings of binary belief

My physicist friend believes that nuclear fusion is a viable energy source. She also believes that her car will stop when she presses the brake pedal. She is willing to bet her life on the latter belief, and in fact does so multiple times every day when she drives to work. She is not willing to bet her life on the former belief. This difference seems like it should be traceable to her differing doxastic attitudes towards the proposition that fusion is viable and the proposition that pressing her brake pedal will stop her car. Yet the Simple Binarist—who is willing to attribute only beliefs, disbeliefs, and suspensions—can make out no difference between my friend's doxastic attitudes towards those propositions. Once the Simple Binarist says my friend believes both propositions, he has said all he has to say.

Now suppose that my physicist friend reads about some new research into nuclear energy. The research reveals new difficulties with tokamak design, which will make fusion power more challenging. After learning of this research, she still believes fusion is a viable energy source. Nevertheless, it seems this evidence should make *some* difference in her attitude towards the proposition that fusion is viable. The Simple Binarist cannot account for this difference; my friend believed the proposition before, and she still believes it now.

What do these two examples show? They don't show that the Simple Binarist embraces any *false* claims—it's *true* that my friend believes the propositions under discussion at the times in question. Instead, they seem to show that the Simple Binarist's descriptive resources aren't fine-grained enough to capture some *further* things we want to say about my friend. Now maybe there's some complicated way the Simple Binarist could account for these examples within his classificatory scheme. Or maybe a complex binarist with more classificatory attitudes in his repetoire than we've given the Simple Binarist could do the trick. But it's most natural to respond to these examples with confidence *comparisons*: my friend is more confident that her brakes will work than she is that fusion is viable, and she is less confident in the viability of fusion after reading the new research than she was before. Even without moving all the way to quantitative degrees of confidence, introducing comparative doxastic attitudes fine-grains our representation in a manner that feels appropriate to the examples.

So far we've discussed difficulties the Simple Binarist will have in *describing* an agent's doxastic attitudes. But along with descriptive adequacy, we often want to work with concepts that allow us to frame plausible *norms*.[3] Historically, epistemologists have often been driven to work with comparative and quantitative doxastic attitudes because of difficulties in framing defensible rational norms for binary belief.

The normative constraints most commonly suggested for binary belief are:

Belief Consistency: Rationality requires the propositions an agent believes to be logically consistent with each other.

Belief Closure: If some of the propositions an agent believes jointly entail a further proposition, rationality requires the agent to believe that further proposition as well.

Belief Consistency and Belief Closure are proposed as necessary conditions for an agent's belief set to be rational. They are also typically proposed as requirements of *theoretical* rather than *practical* rationality.

Practical rationality concerns connections between an attitude and action. Our earlier contrast between my friend's fusion beliefs and her braking beliefs was a practical one; it concerned how those doxastic attitudes influenced her betting *behavior*. Our other problematic example for the Simple Binarist was a purely theoretical one, having to do with my friend's fusion beliefs as evidence-responsive representations of the world (and without considering those beliefs' consequences for her acts).

What kinds of constraints does practical rationality place on attitudes? In Chapter 7 we'll see that if an agent's preferences fail to satisfy certain axioms, this can lead to a disastrous course of actions known as a "money pump". Practical rationality therefore requires agents' preferences to satisfy those axioms. Similarly, we'll see in Chapter 9 that if an agent's credences fail to satisfy the probability axioms, her betting behavior is susceptible to a troublesome "Dutch Book". This fact has been used to argue that practical rationality requires credences to satisfy the probability axioms.

One might think that practical rationality provides all the rational constraints there are.[4] The standard response to this proposal invokes Pascal's Wager. Pascal (1670/1910, Section III) argues that it is rational to believe in the existence of the Christian god because if that belief is true, having believed will yield vast benefits in the afterlife. On the other hand, if the belief is false whether one believed it or not won't have nearly as dramatic consequences. Assuming Pascal has gauged the consequences right, they seem to provide some sort of reason for maintaining religious beliefs. Nevertheless, if an agent's evidence points much more strongly to atheism than to the existence of a deity, it feels like there's a sense of rationality in which religious belief would be a mistake. This is **theoretical rationality**, a standard that assesses representational attitudes considered as such, rather than considering how they influence action. Belief Consistency and Closure are usually offered as requirements of theoretical rationality. The idea is that a set of beliefs has failed as a responsible representation of the world if it contradicts itself or fails to admit its own logical consequences.[5]

The versions of Belief Consistency and Closure I've stated above are pretty implausible as genuine rational requirements. Belief Closure, for instance, requires an agent to believe any arbitrarily complex proposition entailed by what she already believes, even if she's never come close to entertaining that proposition. And since any set of beliefs has infinitely many logical consequences, Closure also requires rational agents to have infinitely many beliefs. Belief Consistency, meanwhile, forbids an agent from maintaining a logically inconsistent set of beliefs even if the inconsistency is so recondite that she is incapable of seeing it. One might find these requirements far too demanding to be rational constraints.

It could be argued, though, that these flaws in Belief Consistency and Closure have to do with the particular way in which I've stated the norms. Perhaps we could make a few tweaks to these principles that would leave their spirit intact while inoculating them against these particular flaws. In Chapter **??** we will consider such tweaks to a parallel set of Bayesian constraints that face similar problems. In the meantime, though, there are counterexamples to Belief Consistency and Closure that require much more than a few tweaks to resolve.

Kyburg (1961) first described the **Lottery Paradox**:

> A fair lottery has sold one million tickets. Because of the poor odds, an agent who has purchased a ticket believes her ticket will not win. She also believes, of each other ticket purchased in the lottery, that *it* will not win. Nevertheless, she believes that at least one purchased ticket *will* win.

The beliefs attributed to the agent in the story seem rational. Yet these beliefs are logically inconsistent—you cannot consistently believe that at least one ticket will win while believing of each ticket that it will lose. So if this set of beliefs is rationally permissible, we have a counterexample to Belief Consistency.

Some defenders of Belief Consistency have suggested that, strictly speaking, it is irrational for the agent in the Lottery to believe her ticket will lose. (If you believe your ticket will lose, why buy it to begin with?[6]) If true, this resolves the counterexample. But it's difficult to resolve Makinson's (1965) **Preface Paradox** in a similar fashion:

> You write a long nonfiction book with many claims in its main text, each of which you believe. In the acknowledgments at the beginning of the book you write, "I'm sure there are mistakes in the main text, for which I take full responsibility."

Many authors write such statements in the prefaces to their books, and it's hard to deny it's rational for them to do so. It's also very plausible that nonfiction authors believe the contents of what they write. Yet if the concession that there are mistakes is an assertion that there is at least one falsehood in the main text, then the belief asserted in the preface is logically inconsistent with belief in all of the claims in the text.[7]

The Lottery and Preface pose a different kind of problem from our earlier examples. The examples with my friend the physicist didn't show that classificatory belief claims were false; they simply suggested that classificatory claims didn't capture all the aspects of doxastic attitudes we would like. The Lottery and Preface, however, are meant to demonstrate that Belief Consistency and Belief Closure—the most natural normative principles for binary belief—are actually false.

An extensive literature has grown up around the Lottery and Preface, attempting to resolve them in a number of ways. One might deny that the sets of beliefs described in the paradoxes are in fact rational. One might find a clever way to establish that those sets of beliefs don't violate Belief Consistency or Belief Closure. One might drop Belief Consistency and/or Belief Closure in favor of alternate normative constraints on binary belief. All of these responses have been tried, and I couldn't hope to adjudicate their successes and failures here.

For our purposes, the crucial point is that while it remains controversial how to square norms for binary belief with the Lottery and Preface, norms for rational credence have no trouble with those examples at all. In Chapter 2 we'll see that Bayesian norms tell a natural, intuitive story about the rational credences to adopt in the Lottery and Preface. The ease with which Bayesianism handles these paradoxes for belief has been seen as a strong advantage for credence-centered epistemology.

1.2 From binary to graded

1.2.1 Comparative confidence

The previous section articulated both descriptive and normative difficulties for restricting one's attention to exclusively classificatory doxastic attitude ascriptions (belief, disbelief, suspension of judgment, etc.). We imagined a Simple Binarist who works only with these kinds of attitudes, and posed various problems for him. The first descriptive problem was that an agent may believe two propositions while nevertheless treating these propositions quite differently when it comes to action. The second descriptive problem was that new evidence may change an agent's doxastic attitude towards a proposition despite the fact that she believes that proposition both before and after incorporating the evidence. We could address both of these shortcomings in a natural fashion by moving beyond strictly classificatory terms and allowing ourselves to make *comparisons* between an agent's levels of confidence in two propositions, or his levels of confidence in a single

Figure 1.1: A partial confidence ordering

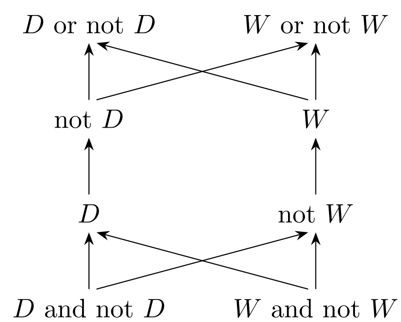

proposition at two different times.

So let's augment the Simple Binarist's resources a bit. We'll still allow ourselves to say that an agent believes, disbelieves, or suspends judgment in a proposition. But let's add the comparative resources to say that an agent is *at least as confident* of one proposition as another, *more confident* in one proposition than another, or *equally confident* of the two. Some of these comparisons follow directly from classificatory claims. For instance, when I say that my friend believes nuclear fusion is a viable energy source, we typically infer that she is more confident in the proposition that fusion is viable than she is in the proposition that fusion is nonviable. But we can also add comparisons that go beyond classificatory information. To the fact that my friend believes both that fusion is viable and that her brakes are functional, we might add that she is *more confident* that her brakes will work than she is that fusion will.

Introducing confidence comparisons between the propositions in a set creates a formal structure called an **ordering** on that set. For example, Figure 1.1 depicts my confidence ordering over a particular set of propositions. Here D represents the proposition that the Democrats will win the next presidential election, and W represents the proposition that anthropogenic global warming has occurred. The arrows indicate *more confident than* relations: for instance, I am more confident that warming either has or hasn't occurred than I am that it has, but I am also more confident that warming has occurred than I am that it has not.

It's important that a confidence ordering may be a **partial ordering**—there may be some pairs of propositions for which the ordering says nothing about the agent's relative confidences. Don't be fooled by the fact that "not D" and "W" are at the same height in Figure 1.1. In that diagram only the *arrows* reflect features of the ordering; the ordering depicted remains silent about whether I am more confident in "not D" or "W". This reflects an important truth about my doxastic attitudes: while I'm more confident in warming than nonwarming and in a Democratic loss than a win, I may genuinely be incapable of making a confidence comparison across those two unrelated issues. In other words, I may view warming propositions and election propositions as *incommensurable*.

We now have the basic elements of a descriptive scheme for attributing comparative doxastic attitudes. How might we add a normative element to this scheme? One popular norm for confidence comparisons is:

Comparative Entailment: For any pair of propositions such that the first entails the second, rationality requires an agent to be at least as confident of the second as the first.

Comparative Entailment is similar in some senses to the Belief Closure norm. Belief Closure says that if a proposition has the high doxastic status of being believed by an agent, any proposition it entails must have that high status as well. Comparative Entailment says that if a proposition receives a particular confidence, any proposition it entails must receive at least that confidence as well.[8] But Comparative Entailment is also plausible on its own terms. It highlights the rational oddness of, say, being more confident that the Yankees are the best baseball team in New York than one is that the Yankees are a baseball team.[9]

Although it's a simple norm, Comparative Entailment has a number of substantive consequences. For instance, assuming we are working with a classical entailment relation on which any proposition entails a tautology and every tautology entails every other, Comparative Entailment requires a rational agent to be equally confident of every tautology and at least as confident of any tautology as she is of anything else. Comparative Entailment also requires a rational agent to be equally confident of every contradiction.

While Comparative Entailment (or something close to it) has generally been endorsed by authors working on comparative confidence relations, there is great disagreement about what additional comparative norms should be accepted. We will present some alternatives in Chapter **??**, when we delve into the technical details of comparative confidence orderings.

1.2.2 Bayesian Epistemology

There is no single view that is Bayesian Epistemology; instead, there are a number of Bayesian epistemologies.[10] Every view I would call a Bayesian epistemology endorses the following two principles:

1. Agents have doxastic attitudes that can usefully be represented by assigning real numbers to claims.

2. Rational requirements on those doxastic attitudes can be represented by mathematical constraints on real numbers closely related to the probability calculus.

The first of these principles is descriptive, while the second is normative—reflecting the fact that Bayesian epistemologies have both desriptive and normative commitments. Most of the rest of this chapter concerns the descriptive element; extensive coverage of Bayesian Epistemology's normative content begins in Chapter 2.[11]

I've articulated the two principles vaguely to make them consistent with the wide variety of views (many of which we'll see later in this book) that call themselves Bayesian epistemologies. For instance, the first principle mentions assigning real numbers to "claims" because some Bayesians use sentences or other entities in place of propositions. Still, the most common Bayesian descriptive approach—and the one we will stick with for most of this book—works with agents' degrees of confidence, measuring the amount of confidence by assigning a real number to a proposition. These degrees of confidence are variously described as "degrees of belief", "graded beliefs" (in contrast with "binary beliefs"), or "credences".[12]

We have already augmented our description of agents' doxastic attitudes by introducing confidence comparisons that go beyond categorical belief/disbelief/suspension terms. What more do we add by moving to a full numerical representation of confidence? Comparative confidence relations introduce orderings—they put things in order. But they cannot tell us how relatively big the gaps are between items in the ordering. Lacking quantitative credal concepts we can say that an agent is more confident in one proposition than she is in another, but we cannot say *how much more* confident she is.

These matters of degree can be very important. Suppose you've been offered a job teaching at a university, but there's another university at which you'd much rather teach. The first university has given you two weeks to respond to their offer, and you know you won't have a hiring decision from the preferred school by then. Trying to decide whether to turn down the offer in hand, you contact a friend at the preferred university. She says you're one of only two candidates for their job, and she's more confident that you'll get the offer than the other candidate. At this point you want to ask *how much more* confident she is in your prospects than the other candidate's. A 51-49 split might not be enough for you to hang in!

Like our earlier brake pedal story, this is an example about the practical consequences of doxastic attitudes. It suggests that distinctions between doxastic attitudes affecting action cannot all be captured by a confidence ordering—important decisions may depend on the sizes of the gaps. Put another way, this example suggests that one needs more than just confidence orderings to do decision theory (which will be the subject of Chapter 7). In Chapter 6, meanwhile, we will use confidence quantities for theoretical purposes: to measure degrees of confirmation. For example, numerical credence values are very important in determining whether a body of experimental evidence supports one scientific hypothesis more than it does another.

These are some of the advantages of numerically measuring degrees of belief. But credal descriptions have disadvantages as well. For instance, numerical representations may offer more specific information than is actually present in the situation being represented. The Beatles were better than the Monkees, but there was no numerical *amount* by which they were better. Similarly, I might be more confident that the Democrats will lose the next election than I am that they will win without there being a fact of the matter about exactly how much more confident I am. Representing my attitudes by assigning precise credence numbers to the proposition that the Democrats will lose and the proposition that they will win attributes a specific confidence gap to me—which may be an *over*-attribution in the actual case.

Numerical degree of belief representations also impose complete commensurability. It is possible to build a Bayesian representation of an agent that does not assign *any* number to a particular proposition, representing the fact that the agent doesn't take any attitude towards that proposition.[13] But once our representation assigns a numerical credence to a particular proposition, that proposition immediately becomes comparable to every other proposition to which a credence is assigned. Suppose I am 60% confident that the Democrats will lose, 40% confident that they will win, and 80% confident that anthropogenic global warming has occurred. One can immediately rank all three of these propositions with respect to my confidence in them. Assigning numerical credences over a set of propositions introduces a **total ordering** on the set, making it impossible to retain any incommensurabilities among the propositions involved (as we wanted to do in Figure 1.1). This is worrying if you think confidence incommensurability is a common and rational feature in real agents' doxastic lives.

Epistemologists sometimes complain that working with numerical credences is unrealistic, because agents "don't have numbers in their heads". This is a bit like refusing to measure gases with a numerical temperature scale because molecules don't fly around with numbers pinned to their backs. The relevant question is whether agents' doxastic attitudes have a level of structure that can be well-represented by numbers, by a partial ordering, by classificatory concepts, or by something else. This is the point at which it's significant to worry whether agents' confidence gaps have important size characteristics, or whether taking attitudes towards any two propositions should automatically make them confidence commensurable. Notice also that there may be no universally best single representation—it may be that different approaches are better in different circumstances. We will return to these issues a number of times in this book.

1.2.3 Relating beliefs and credences

I've said a lot about representing agents as having various doxastic attitudes. But presumably these attitudes aren't just things we can *represent* agents as having; presumably agents actually *have* at least some of the attitudes in question. The metaphysics and ontology of doxastic attitudes raise a huge number of questions. For instance: What *is* it—if anything—for an agent to genuinely possess a mental attitude beyond being usefully representable as having such? Or: If an agent can have both binary beliefs and degrees of belief in the same set of propositions, how are those different sorts of doxastic attitudes related? The latter question has generated a great deal of discussion, which I cannot hope to summarize here. Yet I do want to mention some of the general issues and best-known proposals.

Suppose some connection is asserted between an agent's beliefs and her degrees of belief. That connection might do any of the following: (1) *define* attitudes of one kind in terms of the other; (2) *reduce* attitudes of one kind to attitudes of the other; (3) assert a *descriptively true* conditional (or biconditional) linking one kind of attitude to the other; (4) offer a *normative constraint* to the effect that any rational agent with an attitude of one kind will have a particular attitude of the other.

For example, the **Lockean thesis** connects believing a proposition with having a degree of confidence in that proposition above a numerical threshold. Taking inspiration from John Locke (1975, Bk. IV, Ch. 15 & 16), Richard Foley entertains the idea that:

> 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 in it. (1993, p. 140)

Foley presents the first sentence—identifying belief with sufficiently high degree of belief—as the Lockean thesis. The latter sentence is presented as following from the former. But notice that the latter sentence's normative claim could be secured by a weaker, purely normative Lockean thesis, asserting only that a rational agent believes a proposition just in case she is sufficiently confident of it.

On any reading of the Lockean thesis, there are going to be questions about exactly how high this threshold must be. One might suggest that the confidence threshold for belief is certainty (i.e. 100% confidence). But many of us believe propositions of which we are not certain, and this seems perfectly rational. Working down the confidence spectrum, it seems that in order to believe a proposition one should be more confident of it than not. But that leaves a lot of space to pin down the threshold between 50% and 100% confidence. Here it may help to suggest that the relevant threshold for belief is vague, or varies with context.

The Lockean thesis also causes problems when we try to layer traditional norms of rational belief and credence on top of it. If we adopt Bayesian probabilistic norms for credence, the Lockean thesis generates rational belief sets for the Lottery and Preface that violate Belief Consistency and Closure. We will see why when we give a probabilistic solution to the Lottery in Chapter 2.

The Lockean thesis works by expressing belief as a particular kind of credence. But we might try connecting these attitudes in the opposite direction. For instance, we might say I have 60% credence that the Democrats will lose the next election just in case I believe that their probability of losing is 60%. The general strategy here is to align my credence in one proposition with belief in a second proposition about the *probability* of the first.

This connective strategy—whether meant definitionally, reductively, normatively, etc.—is viewed nowadays as unlikely to succeed. For one thing, it requires thinking that whenever a (rational) agent has a degree of confidence, she also has a belief about probabilities. David Christensen (2004, Ch. 2) wonders what these beliefs about probability are supposed to be beliefs about. In Chapter 5 we will explore various "interpretations of probability" that attempt to explain the meaning of probability claims. The details need not concern us here; what matters is that for each possible interpretation, it's implausible to think that whenever a (rational) agent has a degree of confidence she (also?) has a belief with that kind of probabilistic content. If probability talk is, for instance, always talk about frequency within a reference class, must I have beliefs about frequencies and reference classes in order to be pessimistic about the Democrats?

Deeper problems also arise with requiring the numerical value of a credence to appear inside a proposition towards which the agent adopts an attitude. We will discuss some of these problems when we cover conditional credences in Chapter 3. Generally, contemporary Bayesians think of the numerical value of a credence not as part of the content towards which the agent adopts the attitude, but instead as an attribute of the attitude itself. I adopt a credence *of 60%* towards the proposition that the Democrats will lose; no proposition *containing* the value 60% is involved.[14]

This is a small sample of the positions and principles that have been proposed relating beliefs to degrees of belief. Further connective principles are available; or one might deny there is any principled connection between the two doxastic categories; or one might deny the existence of the attitudes in one category altogether. Going forward, we will assume that it can at least be useful to represent agents as having numerical degrees of belief. We will touch on binary beliefs only rarely.

1.3 The rest of this book

Hopefully I have now given you some sense of what credences are, and of why one might incorporate them into one's epistemology. Our first task in Chapter 2 will be to develop a Bayesian formalism in which credences can be descriptively represented. After that, much of our focus will be on the norms Bayesians require of rational degrees of belief.

There is a great deal of disagreement among Bayesians about exactly what these norms should be. Nevertheless, we can identify five core normative Bayesian rules: Kolmogorov's three probability axioms for unconditional credence, the Ratio Formula for conditional credence, and Conditionalization for updating credences over time. These are not core rules in the sense that all Bayesian epistemologists agree with them. Some Bayesians accept all five rules and want to add more; some don't even accept these five. They are core in the sense that one needs to understand them in order to understand any further Bayesian position proposed.

Part II of this book is primarily concerned with the five core Bayesian rules. Chapter 2 covers Kolmogorov's axioms; Chapter 3 covers the Ratio Formula; and Chapter 4 covers Conditionalization. Chapter 5 then discusses a variety of norms Bayesians have proposed either to supplement or to replace the core five.

The presence of all these alternatives raises the question of why we should accept any of these rules as genuinely normative to begin with. To my mind, one can see the advantages of Bayesianism best by seeing its consequences for applications. For instance, I've already mentioned that Bayesian credence norms match nicely with a natural story about doxastic attitudes in the Lottery Paradox. Part III of the book discusses the two historically most important applications of Bayesian Epistemology: confirmation theory (Chapter 6) and decision theory (Chapter 7).

Along with their benefits in application, Bayesian normative rules have been directly defended with a variety of philosophical arguments. I discuss the three most popular arguments in Part IV, and explain why I find each ultimately unconvincing. Chapter 8 discusses Representation Theorem Arguments; Chapter 9 Dutch Books; and Chapter 10 arguments based on the goal of accurate credences.

Finally, a number of important challenges have been raised to Bayesian Epistemology—both to its descriptive framework and to its normative rules. Many of these (though admittedly not all) are covered in Part V.

1.4 Exercises

Problem 1.1. Explain why (given a classical logical entailment relation) Comparative Entailment requires a rational agent to be equally confident of every contradiction.

Problem 1.2. What do *you* think the agent in the Lottery Paradox should believe? In particular, should she believe of each ticket in the lottery that that ticket will lose? Does it make a difference how many tickets there are in the lottery? Explain and defend your answers.

Problem 1.3. Suppose we have a confidence ordering consisting of *only* the relations depicted by arrows in Figure 1.1. So, for example, the agent in question is more confident in "W" than "not W" (because there's an arrow from the latter to the former), but is not more confident in "W" than "W and not W" (because there is no arrow connecting the two).

(a) Explain why this ordering does not satisfy Comparative Entailment.

(b) Describe all the arrows that would have to be added to the diagram to make the ordering satisfy Comparative Entailment.

Problem 1.4. Assign numerical confidence values (between 0% and 100%, inclusive) to each of the propositions mentioned in Figure 1.1. These confidence values should be arranged so that if there's an arrow in Figure 1.1 from one proposition to another, then the first proposition has a lower confidence value than the second.

1.5 Further reading

Classic Texts

Henry E. Kyburg Jr (1970). Conjunctivitis. In: *Induction, Acceptance, and Rational Belief.* Ed. by M. Swain. Boston: Reidel, pp. 55–82

David C. Makinson (1965). The Paradox of the Preface. *Analysis* 25, pp. 205–7

Classic discussions of the Lottery and Preface Paradoxes (respectively), by the authors who introduced these paradoxes to the philosophical literature.

Extended Discussion

Richard Foley (1993). *Working Without a Net.* Oxford: Oxford University Press

David Christensen (2004). *Putting Logic in its Place.* Oxford: Oxford University Press

Foley and Christensen each discuss the relation of binary beliefs to graded, and the troubles for binary rationality norms generated by the Lottery and Preface Paradoxes. They end up leaning in different directions: Christensen stresses the centrality of credence to norms of theoretical rationality, while Foley emphasizes the role of binary belief in a robust epistemology.

Notes

[1]While Bayesian Epistemology has focused mostly on doxastic representational attitudes, recent years have seen attemps to apply Bayesian ideas to the study of knowledge (see, for instance, (Moss 2013)). But so far no degree-theoretic approach to knowledge is nearly as systematic, well-worked out, or generally-acknowledged as the Bayesian theory of probabilistic credence.

[2]John Bengson, who has greatly helped me with this chapter, brought up the interesting historical example of how we might characterize David Hume's (1739–40/1978) theory of belief vivacity in classificatory/comparative/quantitative terms.

[3]On some epistemologies the descriptive and normative projects cannot be prized apart, because various normative conditions are either definitional or constitutive of what it *is* to possess particular doxastic attitudes. See, for instance, (Davidson 1984) and (Kim 1988).

[4]See, for example, (Kornblith 1993). Kornblith has a response to the Pascalian argument I'm about to offer, but chasing down his line would take us too far afield.

[5]Has Pascal demonstrated that practical rationality requires religious belief? I defined practical rationality as concerning an attitude's connection to action. One odd aspect of Pascal's Wager is that it seems to treat believing as a kind of action in itself. Yet many have wondered whether we have the kind of control over our beliefs that would allow us to deliberately put Pascal's arguments into practice—even if we found them persuasive.

Still, the crucial point is that the pressure to honor our atheistic evidence doesn't seem immediately connected to action in any way. This establishes a standard of theoretical rationality distinct from practically rational concerns.

[6]This is why I never play the lottery.

[7]If you find the Preface Paradox somehow unrealistic or too distant from your life, consider that (1) you have a large number of beliefs (each of which, presumably, you believe); and (2) you may also believe (quite reasonably) that at least one of your beliefs is false. This combination is logically inconsistent.

[8]Comparative Entailment also shares the two flaws we pointed out for Belief Consistency and Closure: (1) as stated, it requires an agent to compare infinitely many ordered pairs of propositions (including propositions the agent has never entertained); and (2) it applies to agents who have not yet realized that an entailment relation holds between a particular pair of propositions.

[9]In an article dated January 2, 2014 on grantland.com, a number of authors made bold predictions for the forthcoming year. Amos Barshad wrote,

> "And so, here goes, my two-part prediction:
>
> 1. The Wu-Tang album will actually come out.
>
> 2. It'll be incredible.
>
> I'm actually, illogically more sure of no. 2."

[10]I.J. Good famously argued in a letter to the editor of *The American Statistician* that there are at least $46,656$ varieties of Bayesians. (Good 1971)

[11]In the philosophy profession these days one sometimes hears discussion of "Formal Epistemology". A formal epistemology is any epistemological theory that uses formal tools. Bayesian Epistemology is just one example of a formal epistemology; other examples include AGM theory (Alchourrón, Gärdenfors, and Makinson 1985) and ranking theory (Spohn 2012).

[12]A very different Bayesian-inspired approach to epistemology assigns propositions numbers obeying the probability calculus, but uses those numbers to represent the propositions' degrees of justification (rather than an agent's confidence in the propositions). I will comment on this degree-of-justification approach in Chapter 6.

[13]I'll mention some details of this move in Chapter XXX.

[14]If we shouldn't think of the number in a numerical credence as part of the content of the proposition towards which the attitude is adopted, how exactly *should* we think of it? I tend to think of the numerical value as a sort of property or adjustable parameter of a particular doxastic attitude-type, credence. An agent adopts a credence towards a specific proposition, and it's a fact about that credence that it has degree 60% (or whatever).

This isn't the only way of thinking about degrees of confidence I'd be willing to accept. We should, however, avoid thinking of degree of confidence as just another property of binary belief, the classificatory concept studied by epistemologists since time immemorial. (This type of thinking might be inspired by the motto "Belief comes in degrees.") Suppose an agent's degree of confidence in some proposition is 1%. You could *say* that the agent has a belief in that proposition; it's just that that particular belief comes with a low degree of confidence. But in this situation I find it hard to maintain that the agent *believes* the proposition in anything like the traditional sense.