Notre Dame Journal of Formal Logic
Volume 47, Number 1, 2006

Book Review

David Christensen. *Putting Logic in its Place: Formal Constraints on Rational Belief*. Oxford University Press, New York, 2004. xii+187 pages.

Christensen distinguishes two kinds of belief, which he calls *binary* and *graded*. Binary belief is qualitative; graded belief is a matter of degree. He argues that the standards of ideal epistemic rationality do not require binary beliefs to be consistent, or closed under logical consequence, but they do require graded beliefs to satisfy the laws of probability. Thus the place of logic—alluded to in the title—is in governing graded, not binary, beliefs.

This is not a technical book; instead it gives very readable discussions of the philosophical issues. Christensen often makes a convincing case that previous discussions have been too facile and for that reason I think anyone interested in epistemic rationality should read this relatively short book. On the other hand, Christensen's arguments for his positive views are less successful; I will show that he has failed to establish either that binary beliefs may violate consistency and closure or that graded beliefs must satisfy the laws of probability.

1 Rationality and Logic

Christensen begins by trying to clarify the concept of rationality that is central to the book. He says that rational beliefs "are those that arise from good thinking, whether or not that thinking was successful in latching on to the truth" (p. 2). He distinguishes two ways in which a belief may be rational: *pragmatic* and *epistemic*. Christensen does not give a general account of this distinction but his examples and discussion suggest this: A belief is pragmatically rational if adopting it is a good means to the believer's ends; it is epistemically rational if it is in accord with the evidence (p. 4). Christensen says his book is concerned with epistemic rationality, not pragmatic rationality.

Since he is concerned with constraints that logic imposes on rational belief, Christensen briefly discusses the nature of logic. He says he is concerned with "formal" logic, but declines to say what that is, except that it is

Printed March 22, 2006
2000 Mathematics Subject Classification: Primary, 03B42, 03B48; Secondary, 60A05
Keywords: degree of belief, preface paradox, Dutch book argument, vague belief

> 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.' (p. 3)

He does not say that logical implications hold in virtue of meanings and he explicitly declines to endorse the claim that "Sulfur is yellow" logically implies "Sulfur is not red" (p. 3). In several other places (pp. 156, 171), Christensen denies that he is assuming a distinction between analytic and synthetic statements. He also says that "there may turn out to be a deep connection between considerations of truth and considerations which make certain claims 'logical' " (p. 2). All this is consistent with the view that logical truths are *a posteriori* propositions (a view associated with Putnam and Quine).

On the other hand, in his Dutch book argument Christensen asserts that bets which will leave an agent worse off in every logically possible situation are "guaranteed" (p. 118) to lose and "the diagnosis can be made *a priori*" (p. 121). This is not consistent with the view that logical truths are *a posteriori* propositions. Thus Christensen's characterization of logic is insufficient to support what he says about logic later in the book.

2 Binary and Graded Belief

In his second chapter, Christensen discusses the nature of binary and graded belief. He argues that both kinds of belief seem to exist and then he considers the possibility of reducing one kind to the other.

2.1 Reduction of binary to graded It has sometimes been claimed that graded beliefs are binary beliefs about probabilities. Christensen has two arguments against this.

2.1.1 First argument Christensen begins thus:

> 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 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. (p. 19)

Christensen supports the last assertion by saying that the agent need have no thought about reference classes or propensities, so that the probabilities in question cannot be those of the frequency or propensity interpretations. He acknowledges that "other objective accounts of probability exist" but says that the two examples he has considered "serve well enough to show how unnatural it is to identify" graded beliefs with binary beliefs about objective probabilities (p. 20). I don't agree, because there are objective probability concepts that are more plausible candidates than the two Christensen considers.

The word "probability" in ordinary language has two senses which I call "physical probability" and "inductive probability." Instead of trying to define these concepts in other terms, I think they are best explained by examples. A suitable example is this: If a coin has either two heads or two tails, then the physical probability that it will land heads when tossed is either 0 or 1, but the inductive probability that it will land heads, given just this evidence, has an intermediate value, plausibly 1/2. The frequency and propensity "interpretations" that Christensen considers are best understood as explicata for the vague ordinary concept of physical probability. That is, they are more precise concepts that are intended to behave similarly to the ordinary concept of physical probability. Since the explicata are more precise than the explicandum, someone who has a belief about a physical probability need not have a belief about either of these explicata. Thus, for all Christensen has said, graded beliefs might be binary beliefs about physical probabilities.

Inductive probability is not the same as subjective probability. For example, when someone says that some hypothesis is probable given some evidence, they are not normally talking about physical probability, but they are also not making a claim about their own psychological state; one can't prove such a probability claim by producing psychological evidence of one's belief in the hypothesis. So inductive probability is another objective concept of probability.

Although Christensen does not consider inductive probability in the argument I am discussing, it plays a crucial role in his book. Consider, for example, the following passage, in which Christensen is explaining what he means by "epistemic rationality":

> 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. 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. (p. 4, italics added)

The terms "likely" and "unlikely" designate a kind of probability, but Christensen does not here mean either subjective or physical probability; his assertions only make sense if understood as being about inductive probability. Similarly, "a belief adopted counter to the evidence" is just a belief that has a low inductive probability given the evidence. I could quote many other passages from Christensen's book that are statements about inductive probability (starting with an example about Ghengis Khan on p. 5).

Someone might think that graded beliefs are binary beliefs about inductive probabilities and I don't see how that would "risk attributing to the agent a belief about matters too far removed from the apparent subject matter of her belief." Certainly it does not require the sorts of beliefs about reference classes or propensities that Christensen used to rule out the frequency and propensity interpretations. So Christensen's first argument is not cogent.

2.1.2 Second argument Christensen next argues:

> 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. (p. 20)

Christensen seems to be reasoning:

(1) Hopes, fears, attractions, and aversions have degrees of strength.

∴ (2) If graded beliefs were reducible to binary beliefs, graded beliefs would be unlike hopes, fears, etc.

But this is invalid, since (2) does not deny that graded beliefs have degrees of strength. Christensen would have a valid argument if he replaced (1) with:

(1′) Hopes, fears, attractions, and aversions are not reducible to anything binary.

However, Christensen has given no reason to believe (1′). I think that having high (low) hope is the same as having a high (low) degree of belief that a favorable event will occur, so to assume that hopes are not reducible to anything binary, in order to argue that degrees of belief are not reducible to binary beliefs, would beg the question.

2.1.3 A better argument The view that graded beliefs are binary beliefs about probabilities can be conclusively refuted by noting that, for any objective concept of probability, people sometimes have a binary belief that the probability of a proposition differs from their degree of belief in that proposition.

For example, if I know that a coin has two heads or two tails, but have no information about which of these is the case, and I know that the coin is about to be tossed, then I have a binary belief that the physical probability of the coin landing heads is 0 or 1, but my degree of belief in this outcome will have an intermediate value, perhaps 1/2. Hence graded beliefs are not binary beliefs about physical probabilities.

Similarly, if I know that E either logically implies H or else logically implies not-H, without knowing which of these is the case, then I have a binary belief that the inductive probability of H given E is either 0 or 1, though my degree of belief in H, given E, will have an intermediate value. Hence graded beliefs are also not binary beliefs about inductive probabilities.

So, Christensen is right that graded beliefs are not binary beliefs about probabilities, even though his arguments for that conclusion are not cogent.

2.2 Reduction of graded to binary After rejecting the reduction of graded belief to binary belief, Christensen says that “a more promising sort of unification would work in the opposite way” (p. 20). The idea here is that a binary belief is a graded belief with a strength that meets a certain threshold. Christensen argues that the threshold needs to be less than certainty and says that it may be vague. He acknowledges that this reduction does not fit perfectly with the way we attribute beliefs, especially in lottery cases, but suggests that the discrepancies might be explained away somehow, “e.g. by invoking principles of conversational implicature” (p. 24). This discussion is in Section 2.2 of Christensen’s book.

In Section 4.4 Christensen discusses this issue again. Here he says that, when his credence in a proposition moves gradually from low to high, or vice versa, “at no point during the process do I seem to experience a discrete qualitative shift in my attitude toward the proposition” (p. 98). He suggests that our attribution of binary beliefs is merely a rough characterization of graded belief as high or not, where what counts as “high” may depend on “contextually determined conversational saliencies” (p. 99). In particular, he is sympathetic to Nozick’s suggestion that “our willingness to attribute belief may depend on what practical matters are at stake” (p. 99). This account differs from the earlier one. It does not try to deal with problem cases by invoking principles of conversational implicature; instead, problems are resolved by allowing that binary beliefs are not a function of graded beliefs alone but depend also on contextual factors.

Christensen supports this later account with an analogy (pp. 96–97, 99–100): We classify dogs as big, small, and medium-sized, but it is plausible that this classification is merely a rough division of the continuous variable of size, where the boundaries may shift depending on context. Christensen thinks it is plausible that our attribution of binary beliefs works similarly; I agree.

3 Consistency and Closure

I now turn to Christensen's claim that ideal epistemic rationality does not require binary beliefs to be consistent or closed under logical consequence. Christensen's argument for this is based largely on the preface paradox, which he presents as follows (pp. 33–34): An author might write a book containing many assertions that the author believes, but also recognize that it is almost inevitable that at least a few of these assertions are mistaken, and say in the preface that she believes the book will be found to contain some errors.

> The problem for deductive consistency is obvious. We naturally attribute to our author the belief, apparently expressed quite plainly in the preface, that the body of her book contains at least one error. We also naturally attribute to her beliefs in each of the propositions she asserts in the body of the book. Every one of these beliefs seems eminently rational. Yet the set of beliefs we have attributed to her is inconsistent. Moreover, the fact that our author, apparently quite reasonably, fails to believe that the body of her book is entirely error-free puts her in violation of the closure requirement. (p. 34)

There are many ways of trying to defend consistency and closure against this argument, and Christensen carefully considers most of them.

One attempted defense is to observe that the beliefs that have been attributed to the author do not really violate either consistency or closure; to get violations of these principles we must suppose in addition that the author has correct beliefs about which claims have and have not been asserted in the book. But Christensen (pp. 37–38) argues, persuasively I think, that we can suppose the author has the required additional beliefs without undermining the force of the example.

Another way to try to defend consistency (but not closure) is to observe that, although it is natural for an author to say "This book undoubtedly contains some errors," it is not natural to make the categorical assertion "This book contains errors." Some have taken this to show that rational authors do not have a binary belief that their book contains errors. But Christensen (pp. 39–44) argues that it can be irrational to fail to believe that one's book contains errors, and again I think his argument is persuasive.

A problem with Christensen's argument is that it ignores his own suggestion (made later in the book) that the binary beliefs a person counts as having depend on the context. If we accept this context-relativity, then the requirement of consistency might be understood in either a strong or a weak form, as follows:

Weak consistency: Epistemic rationality requires that, for every context, the binary beliefs that a person has in that context are logically consistent.

Strong consistency: Epistemic rationality requires that all the binary beliefs that a person has in some context or other are logically consistent.

When an author makes a statement in a preface about errors in the book, the context is arguably not the same as when the specific claims in the book are being argued. Christensen himself agrees that the former "oversteps, in some intuitive sense, the context of inquiry" (p. 93). Hence the Preface Paradox is only an objection to strong consistency, not weak consistency.

The situation with closure is similar. This requirement might be understood in either of the following ways:

Weak closure: Epistemic rationality requires that, for every context, the binary beliefs that a person has in that context are closed under logical consequence.

Strong closure: Epistemic rationality requires that all the binary beliefs that a person has in some context or other are closed under logical consequence.

Christensen argues that closure is not a correct requirement because it can be irrational to believe that all the assertions in a book are true. Since a context in which all the assertions in the book are being considered is arguably different to contexts in which only one specific assertion is being considered, this is at best an argument against strong closure.

Defenders of consistency and closure, who accept that binary belief is context-relative, can say that what they mean to defend is weak consistency and weak closure. The Preface Paradox can then be dealt with by saying the following: The various claims that the author defends in the body of the book are believed in the context in which they are asserted, where the issue may be simply whether or not one particular claim is correct. But in the preface, the issue is whether all these claims are correct, and in that context the author does not believe all those claims. Thus our rational author does not violate either weak consistency or weak closure.

I think that weak consistency and weak closure are plausible conditions. I can even suggest a reason why they are true; namely, unless a person is confused (irrational), we adjust the thresholds in each context so as not to attribute beliefs that violate weak consistency or weak closure.

So, Christensen's arguments against consistency and closure are not cogent objections to the plausible forms of those principles. This has happened because he did not take seriously his own suggestion that "our belief-attributing practice" is "sensitive to contextually determined conversational saliencies" (p. 99).

4 Probabilism

Christensen claims that degrees of belief that are epistemically rational and have precise numerical values must satisfy the elementary laws of probability. He presents two arguments for this view, a Dutch book argument and a representation theorem argument.

Christensen is dissatisfied with other people's versions of both arguments for two reasons. First, they postulate a definitional connection between degree of belief and preference, whereas Christensen believes these are separate things that are only linked normatively. Second, these arguments consist in showing that violations of the laws of probability lead to irrational preferences, but Christensen thinks that any such argument "invites the suspicion that it is addressed to pragmatic, not epistemic, rationality" (p. 140). Christensen tries to restate the Dutch book and representation theorem arguments in ways that avoid these sources of dissatisfaction.

4.1 Dutch book argument In [4] I criticized an earlier Dutch book argument by Christensen for using unclear concepts and implausible premises. The argument that Christensen presents in this book avoids the specific objections I made to his earlier argument, but it still uses unclear concepts and implausible premises, as I will now show.

Christensen begins with the following explanation:

> [L]et us say that an agent's degree of belief in a certain proposition sanctions a bet as fair if it provides justification for evaluating the bet as fair—i.e. for being indifferent to taking either side of the bet. Clearly, this connection depends in any given case on the agent's values. If an agent values roast ducks more than boiled turnips, her belief that a coin is unbiased will not sanction as fair a bet in which she risks a roast duck for a chance of gaining a boiled turnip on the next coin flip. If she values the two equally, however, and values nothing else relevant in the context, she should be indifferent to taking either side of a bet, at one duck to one turnip, on the next flip of a coin that she believes to be fair. (p. 116)

Here and elsewhere, Christensen talks as if sanctions-as-fair is a binary relation between a degree of belief and a bet, but he also says that it depends on the agent's values. So I take it that sanctions-as-fair is really a ternary relation between a degree of belief, a set of values, and a bet. Christensen's definition may then be expressed as follows.

Definition 4.1 A degree of belief and a set of values *sanction a bet as fair* if having this degree of belief and these values would justify an agent in being indifferent between the two sides of the bet.

Christensen's argument also uses the concept of a *simple agent*. This is an imaginary person who has linear utility for money and who values nothing other than money. Christensen's first premise is

> *Sanctioning.* A simple agent's degrees of belief sanction as fair monetary bets at odds matching his degrees of belief. (p. 117)

In a footnote, Christensen explains that "if one's degree of belief in a proposition P is r, the matching odds would be $\$r : \$(1-r)$." From this and Definition 4.1 it follows that what Sanctioning asserts is the following.

Proposition 4.2 *A simple agent who has degree of belief r in proposition P is justified in being indifferent between the two sides of a bet on P at odds of $r : 1-r$.*

Let P be that a certain coin will land heads on the next toss, and let α be a simple agent whose only evidence relevant to P is that the coin was purchased in a magic shop and has landed tails on every one of the past 100 tosses. Suppose further that, despite this evidence, α has a degree of belief 1/2 in P. According to Proposition 4.2, α is justified in being indifferent between the two sides of a bet on P at even odds. But this is not so; on the contrary, α ought to prefer betting against P to betting on it. Hence Proposition 4.2 is false.

Christensen could avoid this objection by replacing Definition 4.1 with something like the following. (The word in boldface is the crucial change.)

Definition 4.3 A degree of belief and a set of values *sanction a bet as fair* if an agent who **justifiably** had this degree of belief and these values would be justified in being indifferent between the two sides of the bet.

With this definition, what Sanctioning asserts is not Proposition 4.2 but rather the following.

Proposition 4.4 *A simple agent who justifiably has degree of belief r in proposition P is justified in being indifferent between the two sides of a bet on P at odds of $r : 1 - r$.*

This is not open to the objection that I made against Proposition 4.2. But now, let α be a simple agent who justifiably believes both P and not-P to degree 1/3. (Christensen, of course, wants to prove that no agent can justifiably have these degrees of belief, but that cannot be *assumed* here without begging the question. Anyway, we can make a counterfactual supposition.) According to Proposition 4.4, α is justified in being indifferent between the two sides of a bet on P at 1:2 odds. However, since α has equal degrees of belief in P and not-P, I think it is more plausible to say that α would be justified in preferring the side with the bigger payoff, that is, in preferring to bet on P rather than against it. This is supported by the fact that α's subjective expected utility is higher for the bet on P than for the bet against P. (The subjective expected utility of a bet is calculated by multiplying each possible outcome of the bet by α's degree of belief that it will happen, then adding the products.) I conclude that Proposition 4.4 is implausible.

I will not give a detailed analysis of the remainder of Christensen's Dutch book argument for the following reasons.

1. I have already said enough to show that the argument is not cogent.
2. The remainder of the argument will have to change if Christensen revises his first premise to meet my objections.
3. As it stands, the remainder of the argument involves multiple obscurities.

I will, however, illustrate what I mean by that last remark.

Christensen's second premise is

> *Bet Defectiveness.* For a simple agent, a set of bets that is logically guaranteed to leave him monetarily worse off is rationally defective. (p. 118)

As words are ordinarily used, a set of bets is not the kind of thing that can be rational or irrational, and hence not the kind of thing that can be "rationally defective," so presumably Christensen is not using this phrase in an ordinary sense. However, Christensen does not say what he means in calling a set of bets "rationally defective."

Christensen's third premise is

> *Belief Defectiveness.* If a simple agent's beliefs sanction as fair each of a set of bets, and that set of bets is rationally defective, then the agent's beliefs are rationally defective. (p. 119)

Christensen does not say whether the defect in the agent's beliefs here is pragmatic or epistemic. If he means to assert that the defect is in epistemic rationality, then there is an equivocation in the two occurrences of "rationally defective," since bets cannot be defective in epistemic rationality; also, there should be some discussion of why the agent's beliefs are defective in epistemic rationality. If he means to assert a defect in pragmatic rationality, then there must be a transition to epistemic rationality later in the argument, but Christensen does not indicate any such transition.

From the premises stated so far and the Dutch Book theorem, Christensen infers

> *Simple Agent Probabilism.* If a simple agent's degrees of belief violate the probability axioms, they are rationally defective. (p. 121)

Christensen then argues that this conclusion can be generalized to all agents. This argument (pp. 122–23) is obscure; as far as I can tell it is only an assertion masquerading as an argument. Also, Christensen at this point repeatedly asserts that degrees of belief are "representations of the world," but a representation of the world must present the world as being a certain way, and degrees of belief do not do that. For example, a degree of belief of 1/2 that it will rain tomorrow does not present the world as raining tomorrow, or as not raining tomorrow, or as partially raining and partially not raining, or as having a physical probability of 1/2 of raining. Perhaps Christensen has a looser conception of "representation," but he does not explain what it is or how it fits into his argument.

I will leave Christensen's Dutch book argument at this point. In view of its false first premise and its obscurities, this argument provides no support for its conclusion.

4.2 Representation theorem argument After presenting his Dutch book argument, Christensen discusses representation theorem arguments. He characterizes representation theorems as having the following form:

> *Representation Theorem.* If an agent's preferences obey constraints C, then they can be represented as resulting from some unique set of utilities U and probabilistically coherent degrees of belief B relative to which they maximize expected utility. (p. 125)

For his representation theorem argument, Christensen adds the following two premises:

> *Preference Consistency.* Ideally rational agents' preferences obey constraints C. (p. 125)

> *Representational Rationality.* If an *ideally rational* agent's preferences can be represented as resulting from unique utilities U and probabilistically coherent degrees of belief B relative to which they maximize expected utility, then the agent's actual utilities are U and her actual degrees of belief are B. (p. 138)

From these assumptions he draws the conclusion:

> *Probabilism.* Ideally rational agents have probabilistically coherent degrees of belief. (p. 125)

Christensen says that this argument "can lend substantial support to probabilism" (p. 139).

Christensen's formulation of Probabilism seems to imply that ideally rational agents have numerically precise degrees of belief. And whether or not that is what Christensen intended Probabilism to say, the premises of his argument do imply that ideally rational agents have numerically precise degrees of belief in all propositions to which the representation theorem applies. Since Christensen says that he wants to leave open the possibility that epistemically rational degrees of belief may be imprecise (p. 149), he should agree that something is wrong with his representation theorem argument.

There is something wrong and it is this: If the conditions C are strong enough to make the Representation Theorem true then they imply that, for any pair of options, an ideally rational agent either prefers one to the other or else is indifferent between them; the agent cannot be undecided. But if C is this strong, then Preference Consistency is implausible; an agent could be ideally rational yet be undecided between some options.

Christensen dismisses objections to Preference Consistency by saying that "Patrick Maher [3] provides very nice explanations of—and defenses against—these objections" (p. 126). I appreciate the compliment, but Christensen is here forgetting that I only defended *some* of the conditions that need to be in C; in particular, I argued that "the connectedness postulate [which says that the agent always either has a preference or is indifferent] is not a requirement of rationality" ([3], p. 20). So I actually argued *against* Preference Consistency.

Christensen could deal with this problem by revising his argument as follows. Let C′ be the conditions needed for a representation theorem *other than connectedness*. Then we have

> *Representation Theorem*$_M$. If an agent's preferences obey constraints C′ and are connected, then they can be represented as resulting from some unique set of utilities U and probabilistically coherent degrees of belief B relative to which they maximize expected utility.

We can replace Preference Consistency by

> *Preference Consistency*$_M$. Ideally rational agents' preferences obey constraints C′.

Representational Rationality can stay as before and the conclusion becomes

> *Probabilism*$_M$. Ideally rational agents who have connected preferences have probabilistically coherent degrees of belief.

Whether this revised version of Christensen's argument is cogent will depend mainly on what is included in the conditions C′. Christensen does not go into details about this and I will not do so either. Instead I want to consider how, even if the argument is sound as far as it goes, it establishes what Christensen is trying to prove.

What Christensen is trying to prove here is that numerically precise epistemically rational degrees of belief satisfy the laws of probability; but that is not what Probabilism$_M$ says. What we can infer from Probabilism$_M$ is that, if an agent has numerically precise degrees of belief that do not satisfy the laws of probability, then that agent either lacks connected preferences or else is not ideally rational. It does not follow that the agent is less than ideally rational. Furthermore, even if the agent has connected preferences, what follows is only that the agent is less than ideally rational *in some way*, not that the agent is *epistemically* irrational. If the term "ideally rational" in Preference Consistency$_M$ referred to epistemic rationality, that premise would be false, since someone can be epistemically rational without having sensible preferences; so "ideally rational" here must include more than epistemic rationality, and hence a failure of it is not necessarily a failure of epistemic rationality.

In a subsequent section Christensen tries to justify drawing a conclusion about epistemic rationality.

> Beliefs are, after all, more than just a basis of action. The defect inherent in beliefs that violate probabilism should be seen as primarily epistemic rather than pragmatic. The epistemic defect shows itself in pragmatic ways, for a fairly simple reason: The normative principles governing preferences must of course take account of the agent's information about how the world is. When the agent's beliefs—which represent that information—are intrinsically defective, the preferences informed by those defective beliefs show themselves intrinsically defective too. But in both cases [the Dutch book and representation theorem arguments], the preference defects are symptomatic of, not constitutive, of the purely epistemic ones. (p. 141)

The passage repeats the mistaken claim that degrees of belief are representational. But putting that aside, I can see no *argument* here that the defect is epistemic, merely an assertion that it is.

I conclude that Christensen's representation theorem argument, like his Dutch book argument, does not support his claim that precise epistemically rational degrees of belief satisfy the laws of probability.

4.3 What should be done Before trying to argue that precise epistemically rational degrees of belief satisfy the laws of probability, one ought to make clear what one means by an "epistemically rational degree of belief."

Christensen says that "a belief adopted counter to the evidence" is not epistemically rational (p. 4). Also, his examples of beliefs that are not epistemically rational are ones that are "unlikely" given the agent's evidence (p. 4) or for which the agent has "no evidence" (p. 5). On the other hand, Christensen declines to take a stand on whether an epistemically rational degree of belief may, or even must, be numerically precise when the evidence does not pick out any unique precise degree of belief as epistemically rational (p. 149).

Perhaps Christensen thinks there is one determinate concept of epistemically rational degree of belief and it is a question for further inquiry whether it allows, or requires, precise degrees of belief when the evidence is vague; but Christensen has not said enough to identify such a concept. My view is that there are various things one could mean by "epistemically rational degree of belief" and Christensen has not settled which of them he is talking about. The following definitions articulate two of the possibilities; they use the concept of inductive probability that I introduced in Section 2.1.1 and they differ only in the boldface words.

Definition 4.5 Let $\mathcal{S}$ be a set of propositions. An agent's degrees of belief in the propositions in $\mathcal{S}$ are *epistemically rational* if and only if, for each $P \in \mathcal{S}$, the agent's degree of belief in P **equals** the inductive probability of P given the agent's evidence.

Definition 4.6 Let $\mathcal{S}$ be a set of propositions. An agent's degrees of belief in the propositions in $\mathcal{S}$ are *epistemically rational* if and only if, for each $P \in \mathcal{S}$, the agent's degree of belief in P **is not outside the range of** the inductive probability of P given the agent's evidence.

According to Definition 4.5, epistemically rational degrees of belief cannot be more or less precise than the corresponding inductive probability. According to Definition 4.6, epistemically rational degrees of belief can be more precise than the corresponding inductive probability but they cannot be less precise.

If Definition 4.6 is what we mean by "epistemically rational degree of belief," and inductive probabilities are sometimes vague, then precise epistemically rational degrees of belief need not satisfy the laws of probability.

Proof Suppose that the inductive probability of P given agent α's evidence is vague. Let r and r' be different numeric values in the range of this inductive probability. *Case (i):* $1 - r'$ is in the range of the inductive probability of not-P given α's evidence. By Definition 4.6, it would be epistemically rational for α to believe P to degree r and not-P to degree $1 - r'$. Since $r' \neq r$, α's degrees of belief violate the laws of probability. *Case (ii):* $1 - r'$ is not in the range of the inductive probability of not-P given α's evidence. Then there must be some number $r'' \neq r'$ such that $1 - r''$ is in the range of the inductive probability of not-P given α's evidence. Then, by Definition 4.6, it would be epistemically rational for α to believe P to degree r' and not-P to degree $1 - r''$. Since $r'' \neq r'$, α's degrees of belief again violate the laws of probability. □

If Definition 4.5 is what we mean by "epistemically rational degree of belief," then precise epistemically rational degrees of belief satisfy the laws of probability if and only if precise inductive probabilities, given any fixed evidence, satisfy the laws of probability. This seems to be the case. Perhaps inductive probabilities are numerically precise only when the evidence specifies relevant physical probabilities, or the alternatives can be partitioned into equally probable alternatives, and in these cases it is uncontroversial that the probabilities satisfy the (elementary) laws of probability.

So we have two definitions of epistemically rational degrees of belief, both of which are consistent with Christensen's characterization of this concept, but depending on which is adopted we get different truth values for the claim that precise epistemically rational degrees of belief satisfy the laws of probability. Hence it is impossible to prove that claim without first giving a more complete account of what is meant by "epistemically rational degree of belief."

I would choose Definition 4.5 and then argue, in the way I did two paragraphs back, that precise epistemically rational degrees of belief satisfy the laws of probability. On this approach, there is no need for a Dutch book or representation theorem argument.

5 Vagueness

5.1 Christensen's constraint As I mentioned earlier, Christensen does not rule out the possibility that some epistemically rational degrees of belief lack precise numeric values and hence do not satisfy the laws of probability. He deals with this by saying:

> Fortunately, it is not difficult to accommodate spread-out beliefs in a way that preserves the intuitive value of the probabilistic model. As various authors have noted, we may represent an agent's belief state not by a single function assigning numbers to propositions, but by a *set* of such functions. The condition of an agent's degree of belief in P being spread out from 0.2 to 0.3 will be represented by her set of belief-functions including members that assign P the numbers from 0.2 to 0.3, but no members assigning P a value outside this range. Instead of requiring that the agent's beliefs be represented by a single probabilistically coherent belief-function, we may require that the agent's beliefs be representable by a set of belief-functions, each of which is probabilistically coherent. (p. 149)

Christensen does not give a more precise statement of the requirement he is here endorsing, but his remarks here suggest the following.

Proposition 5.1 *If an agent's degrees of belief are epistemically rational then there exists a set W of probability functions such that, for each proposition P, the agent's degree of belief in P is spread out over $\{p(P) : p \in W\}$.*

For this to be true, epistemically rational degrees of belief must always be spread out over some definite set of numbers, and that seems obviously wrong. In a footnote attached to the passage I just quoted, Christensen acknowledges this problem and responds to it as follows:

> Now it seems to me fairly plausible that, if rational attitudes toward propositions may be spread out along ranges of degrees of confidence, those ranges will themselves have vague boundaries—there may well be some vagueness in which precise degrees of belief the evidence rules out. But this point is compatible with the ranges' providing a vastly improved model of spread-out belief. Consider an analogy: We might represent Lake Champlain as stretching from latitude 43:32:15 N to 45:3:24 N. We would realize, of course, that there really aren't non-vague southernmost and northernmost points to the lake; lakes are objects that lack non-vague boundaries. But representing the lake as ranging between these two latitudes is sufficiently accurate, and vastly better than representing the lake's location by picking some single latitude in between them. Similarly, we might represent an ideally rational agent's attitude toward P in a certain evidential situation as ranging from 0.2 to 0.3. We may well do this while realizing that the lowermost and uppermost bounds on degrees of confidence allowed by the evidential situation are vague. But this representation may yet be very accurate, and a considerable improvement over representing the agent's attitude by a single degree of belief. (Thanks to Mark Kaplan for help on this point.) (pp. 149n–150n)

Christensen says here that representing a vague degree of belief as ranging over an interval "may yet be very accurate," but he offers no reason to think that such a representation *is in fact* very accurate. Furthermore, consideration of actual cases suggests that it isn't very accurate, at least if epistemically rational degrees of belief can be as vague as the corresponding inductive probabilities. For example, the inductive probability that there is intelligent life in the Andromeda galaxy, given my evidence or yours, has a relatively large indeterminacy in its boundaries, and the same is true in many other cases. The boundaries of epistemically rational degrees of belief are not as definite as the boundaries of Lake Champlain.

In this footnote, Christensen does not say what constraint on rational degrees of belief he is now endorsing. His overall claim in this section is that when "credences in certain propositions are spread out, they are still constrained by coherence" (p. 150), so he must endorse some constraint or other. One might guess from what he says that he means to endorse the following modification of Proposition 5.1 (the only change is the boldfaced word).

Proposition 5.2 *If an agent's degrees of belief are epistemically rational then there exists a set W of probability functions such that, for each proposition P, the agent's degree of belief in P is* **approximately** *spread out over* $\{p(P) : p \in W\}$.

This is not much of a constraint in the absence of any specification of what counts as being "approximately spread out over $\{p(P) : p \in W\}$," though any such specification seems arbitrary. And even if we leave it vague, there seems no good reason to believe that Proposition 5.2 is true, and the considerations I raised in the preceding paragraph suggest it is false.

5.2 Kaplan's constraint Christensen's treatment of vague degrees of belief is brief and he says in a footnote:

> See Kaplan ([2], ch. 1, sect. V) both for an extended argument in support of using sets of probability assignments to represent rational epistemic states, and for references to various implementations of this strategy.

So I will now examine Kaplan's "implementation of this strategy," after which I will examine Kaplan's "extended argument."

Kaplan argues that rational simple agents satisfy a condition he calls "Modest Probabilism" ([2], p. 21), and he later asserts (pp. 43–44) that more sophisticated arguments by other authors show that the restriction to simple agents is dispensable. He also argues (pp. 40–43) that his claims apply to epistemic rationality. It follows that Kaplan endorses Proposition 5.3.

Proposition 5.3 *If your degrees of belief are epistemically rational then there exists a nonempty set W of probability functions such that*

(i) *you believe P to the same degree that you believe Q iff $p(P) = p(Q)$ for all $p \in W$; and*
(ii) *you believe P to a higher degree than Q iff $p(P) \geq p(Q)$ for all $p \in W$ and $p(P) > p(Q)$ for some $p \in W$.*

A weaker claim implied by Proposition 5.3 is the following.

Proposition 5.4 *If your degrees of belief are epistemically rational then there exists a nonempty set W of probability functions, and a property φ of sets of ordered pairs of numbers, such that you believe P to a higher degree than Q iff $\varphi\{\langle p(P), p(Q)\rangle : p \in W\}$.*

I will now show Proposition 5.4 is false. I assume here a conception of epistemic rationality according to which epistemically rational degrees of belief may be as vague as the corresponding inductive probabilities.

Let P be that there is intelligent life in the Andromeda galaxy. Suppose a ball is to be drawn randomly from an urn containing balls numbered from 1 to 1000, and let Q_n be the proposition that the ball drawn has a number less than or equal to n. My degree of belief in P is higher than my degree of belief in Q_0 and less than my degree of belief in Q_{1000}. However, since my degree of belief in P is vague, there is no n such that my degree of belief in P is higher than my degree of belief in Q_n but not higher than my degree of belief in Q_{n+1}. From these facts and Proposition 5.4 it follows that my degrees of belief are epistemically irrational.

Proof Suppose there is a set W that satisfies the conditions of Proposition 5.4 with respect to me. Let N be the set of those n such that

$$\varphi\{\langle p(P), p(Q_n)\rangle : p \in W\}.$$

By Proposition 5.4 and the facts about me cited above, $0 \in N$ and $1000 \notin N$. Hence there exists an n^*, $0 \leq n^* \leq 999$, such that $n^* \in N$ and $n^* + 1 \notin N$. By Proposition 5.4, it follows that I believe P to a higher degree than Q_{n^*} but not to a higher degree than Q_{n^*+1}. Since there is no n^* for which this is true, the assumption from which I started is false, that is, there is no set W which satisfies the conditions of Proposition 5.4 with respect to me. So, by Proposition 5.4, my degrees of belief are epistemically irrational. □

I don't deny that this conclusion is true, but it doesn't follow from the facts about my degrees of belief that I reported. Therefore, Proposition 5.4 is false. It follows that Proposition 5.3 is false, and so is any other proposition that implies Proposition 5.4.

5.3 Kaplan's argument Christensen said to "see Kaplan ([2], ch. 1, sect. V) . . . for an extended argument in support of using sets of probability assignments to represent rational epistemic states," so I will now consider Kaplan's argument.

Sections I – IV of chapter 1 of Kaplan's book argue that Proposition 5.3 is true of simple agents, and in Section VII (pp. 43–44) he says that more elaborate arguments by other authors show that the restriction to simple agents is dispensable. In Section V Kaplan argues that the set W in Proposition 5.3 sometimes ought to have more than one member, but this assumes the correctness of Proposition 5.3, which is what we are concerned with here. So the relevant argument is not in the section that Christensen cited. I will therefore focus on Kaplan's argument, in Sections I–IV of his book, that Proposition 5.3 is true of simple agents.

Kaplan uses the following terminology:

> **Definition.** A is a *well-mannered state of affairs* just in case, for some set of mutually exclusive and jointly exhaustive hypotheses, $\{P_1, \ldots, P_n\}$, and some set of real numbers $\{a_1, \ldots, a_n\}$, A is identical to (\$$a_1$ if $P_1, \ldots,$ \$$a_n$ if P_n). (p. 4)

One of Kaplan's premises is the following.

Proposition 5.5 *If you are rational simple agent then there exists a nonempty set V of assignments of monetary values to well-mannered state of affairs such that*

(i) *you are indifferent between A and B iff $\upsilon(A) = \upsilon(B)$ for all $\upsilon \in V$; and*
(ii) *you prefer A to B iff $\upsilon(A) \geq \upsilon(B)$ for all $\upsilon \in V$ and $\upsilon(A) > \upsilon(B)$ for some $\upsilon \in V$.*

Kaplan asserts this on pp. 11–12 and it is implied by the "Modest Connectedness" principle that he states on p. 13. A weaker claim, implied by Proposition 5.5, is the following.

Proposition 5.6 *If you are rational simple agent then there exists a nonempty set V of assignments of monetary values to well-mannered states of affairs, and a property φ of sets of ordered pairs of numbers, such that you prefer A to B iff $\varphi\{\langle \upsilon(A), \upsilon(B)\rangle : \upsilon \in V\}$.*

Proposition 5.6 is false for essentially the same reason that Proposition 5.4 is false. To see this, let P and Q_n be as before, let A give \$1 if P is true and nothing otherwise, and let B_n give \$1 if Q_n is true and nothing otherwise. I prefer A to B_0 and I prefer B_{1000} to A. However, since my degree of belief in P is vague, there is no n such that I prefer A to B_n but do not prefer A to B_{n+1}. It follows from Proposition 5.6 that I am not a rational simple agent. (The proof is parallel to the proof following Proposition 5.4.) But while it is true that I am not a rational simple agent, the facts that I just reported are not inconsistent with my being a rational simple agent. Therefore, Proposition 5.6 is false. Hence the stronger Proposition 5.5 that Kaplan assumes is also false, and so Kaplan does not have a sound argument for the claim that rational degrees of belief are representable by sets of probability functions.

5.4 Explication Suppose we agree—as most people today do—that the following is false.

Proposition 5.7 *If a person's degrees of belief are epistemically rational then there exists a probability function that represents the person's degrees of belief.*

The question is how to modify this to get something true. The usual approach, endorsed by Christensen and Kaplan (and also formerly by me), takes the appropriate modification to be the following.

Proposition 5.8 *If a person's degrees of belief are epistemically rational then there exists a* **set of probability functions** *that represents the person's degrees of belief.*

But I have argued that this is also false. Rather, I now think that the appropriate modification is this.

Proposition 5.9 *If a person's degrees of belief are epistemically rational then there exists a probability function that* **is a good explicatum for** *the person's degrees of belief.*

I use the term "explicatum" in Carnap's sense [1, §§2,3]. Thus a good explicatum is a precise concept which is similar to a given vague concept (called the "explicandum") in ways that matter for the purposes we have in view and which is simple and theoretically fruitful. The activity of proposing an explicatum for a given explicandum is called "explication."

When Christensen tries to deal with the problem of vague boundaries of vague degrees of belief, what he claims is that representing these degrees of belief by a set of numbers "may yet be very accurate" (p. 150n). By contrast, a good explicatum need not be a very accurate representation of its explicandum. As Carnap says:

> The explicatum is to be *similar to the explicandum* in such a way that, in most cases in which the explicandum has so far been used, the explicatum can be used; however, close similarity is not required, and considerable differences are permitted. ([1], p. 7)

The differences are permitted because we want the explicatum to be precise, simple, and fruitful, whereas the explicandum may be none of these things.

Christensen also claims that representing degrees of belief by a set of numbers may be "a considerable improvement over representing the agent's attitude by a single degree of belief" (p. 150n). The only reason for this that is suggested by his discussion is that it may be a more accurate representation. However, even if that were true (which I doubt), it would not follow that sets of numbers provide a better explicatum than a single number, since use of a single number is simpler, and likely to be more fruitful, than using a set of numbers.

I can't prove that Proposition 5.9 is true, but I can point to two facts in its favor. First, unlike Proposition 5.8, Proposition 5.9 has not been refuted. Second, probability functions that are good explicata for epistemically rational degrees of belief have been defined for some simple situations of real interest; for an example, see [5]. Further constructive work of this kind is much to be desired.

6 Human Limitations

Christensen's discussion of vagueness occupies a relatively small part of his last chapter. Most of that chapter is concerned with the objection that he has proposed standards of epistemic rationality that cannot be met by human beings. I agree with what Christensen says about this objection, but I think that the objection seems pertinent only because Christensen is not sufficiently explicit about what he means by "epistemic rationality." For example, if he were to adopt Definition 4.5, or even Definition 4.6, it would be clear that he is really talking about the properties of inductive probabilities, and there is no plausible reason to think that humans must be capable of having all their degrees of belief match those.

7 Conclusion

In this review I have made three main criticisms of Christensen's book:

1. When Christensen argues that binary beliefs need not be consistent or closed under logical consequence, his argument is vitiated by a failure to take account of his own suggestion that binary belief is a function of context.
2. Christensen's Dutch book and representation theorem arguments are deeply flawed and provide no support for the view that precise graded beliefs should satisfy the laws of probability. Furthermore, Christensen could not cogently argue for that conclusion without first specifying more carefully what he means by an epistemically rational degree of belief.
3. Christensen's attempt to extend the laws of probability to vague degrees of belief fails for a reason that he himself recognizes, and his attempt to deal with that is poorly articulated and supported by no argument. Christensen refers to Kaplan for further articulation and an argument, but Kaplan's argument is unsound and the position he defends is false.

In addition to making these criticisms, I have indicated what I think is a better way to approach each of these three issues.

References

[1] Carnap, R., *Logical Foundations of Probability*, The University of Chicago Press, Chicago, 1950. 2d edition 1962. Zbl 0040.07001. MR 0040253. 148

[2] Kaplan, M., *Decision Theory as Philosophy*, Cambridge University Press, Cambridge, 1996. Zbl 0885.62004. MR 1686302. 145, 146

[3] Maher, P., *Betting on Theories*, Cambridge Studies in Probability, Induction, and Decision Theory. Cambridge University Press, Cambridge, 1993. MR 1297541. 142

[4] Maher, P., "Depragmatized Dutch book arguments," *Philosophy of Science*, vol. 64 (1997), pp. 291–305. MR 1602229. 139

[5] Maher, P., "Probability captures the logic of scientific confirmation," pp. 69–93 in *Contemporary Debates in Philosophy of Science*, edited by C. R. Hitchcock, Blackwell, Oxford, 2004. 148

Acknowledgments

Thanks to David Christensen for comments.

Patrick Maher
Department of Philosophy
University of Illinois at Urbana-Champaign
Urbana IL 61801
patrick@maher1.net
http://patrick.maher1.net