Australasian Journal of Philosophy

ISSN: 0004-8402 (Print) 1471-6828 (Online) Journal homepage: http://www.tandfonline.com/loi/rajp20

Belief, Credence, and the Preface Paradox

Alex Worsnip

To cite this article: Alex Worsnip (2016) Belief, Credence, and the Preface Paradox, Australasian Journal of Philosophy, 94:3, 549-562, DOI: 10.1080/00048402.2015.1084343

To link to this article: http://dx.doi.org/10.1080/00048402.2015.1084343

AUSTRALASIAN JOURNAL OF PHILOSOPHY, 2016
VOL. 94, NO. 3, 549–562
http://dx.doi.org/10.1080/00048402.2015.1084343

Belief, Credence, and the Preface Paradox

Alex Worsnip

New York University

ABSTRACT

Many discussions of the 'preface paradox' assume that it is more troubling for deductive closure constraints on rational belief if outright belief is reducible to credence. I show that this is an error: we can generate the problem without assuming such reducibility. All that we need are some very weak normative assumptions about *rational* relationships between belief and credence. The only view that escapes my way of formulating the problem for the deductive closure constraint is in fact itself a reductive view: namely, the view that outright belief is credence 1. However, I argue that this view is unsustainable. Moreover, my version of the problem turns on no particular theory of evidence or evidential probability, and so cannot be avoided by adopting some revisionary such theory. In sum, deductive closure is in more serious, and more general, trouble than some have thought.

KEYWORDS deductive closure; coherence requirements; rational belief; credence; preface paradox

1. Introduction

According to a tempting line of thought, a perfectly rational doxastic subject's belief set will be deductively closed under entailment. That is to say, the following requirement of rationality holds:

> **Deductive Closure Constraint.** Necessarily, rationality requires of S that if S's total belief set logically entails some proposition p then S believes p.[1]

Various worries might be raised about this constraint's being too demanding.[2] Still, even those who are live to such worries might maintain that obeying the Deductive Closure Constraint is at least an *ideal*.[3] Or they might maintain that having beliefs that are closed under logical entailment is at least *permitted*:

> **Deductive Closure Permission.** Necessarily, if S's current total belief set is rationally permissible and S's total belief set logically entails some proposition p, then it is rationally permissible for S to believe p (while also retaining all of the beliefs in her current total belief set).[4]

[1] In this paper, I use the lower-case 'p' for propositions and the upper-case 'P' for sentences that pick out those propositions.
[2] See, e.g., Harman [1986].
[3] See, e.g., Kaplan [1996: 113].
[4] This is a more general version of what sometimes gets called the 'conjunction rule' (cf. Stalnaker [1984: 93] and Foley [1993: 166]), on which, if it's rationally permissible to (believe p and believe q), then it's rationally permissible to believe (p & q).

This captures the rough idea that it is permissible to extend one's belief set by performing logical operations on one's existing beliefs.

'Preface paradox' cases, however, question even whether these weaker claims are true. They purport to show that a subject can be *rationally required* to withhold belief from (or even believe the negation of) a recognized logical consequence of her other beliefs, even if those other beliefs are themselves rational. Thus, they suggest that the Deductive Closure Constraint is not even an ideal, and that the Deductive Closure Permission is false.

2. The Original Preface Cases

In the classical preface paradox case, an author has written a long book containing a large number of claims—yet indicates in the preface of the book that the book no doubt contains errors.[5] This conjunction of claims—those in the body of the book plus that in the preface—appears to be inconsistent. Yet it seems that, even if every claim in the book is responsible and rational when taken by itself, the author is right to think that there are probably errors—and to withhold belief from the claim that the book contains no errors. Moreover, to believe that the book contains no errors seems irrational, since even though each individual claim has only a small chance of being false, these risks of error aggregate to make it very likely that there is at least one error in the book.

Unfortunately, although this presentation of the challenge to the Deductive Closure Constraint from risk aggregation is nicely concrete and vivid, the details of the case introduce distracting features that are in fact immaterial to the central point being made. First, as a few philosophers have pointed out,[6] strictly speaking the claims in the book plus the preface claim do *not* actually form an inconsistent set. Suppose that, in the actual world, the author writes some particular book consisting of the set of propositions $\{p_1, p_2, \ldots, p_n\}$. Call the proposition that there is at least one error in the author's book *e*. Now consider some other possible world in which the set $\{p_1, p_2, \ldots, p_n\}$ contains only true propositions, but in which the author writes a slightly different book, identical to the book she writes at the actual world except for containing $\neg p_2$ instead of p_2. That is a world in which the set $\{p_1, p_2, \ldots, p_n, e\}$ contains only true propositions, and yet where e is true. So, the set $\{p_1, p_2, \ldots, p_n, e\}$ is consistent. In order to make it inconsistent, we have to add to this set the proposition that the author's book contains the propositions $\{p_1, p_2, \ldots, p_n\}$. That does the trick of making the set inconsistent. But it is not plausible that the author has a correct single belief about exactly which propositions her book contains, or that she is capable of forming such a belief.

Second, one may doubt whether the writers of books really do believe every proposition that their books contain.[7] Belief-ascription is a tricky matter, and the mere fact that someone wrote a book saying something does not establish that she (outright) believes the content of what is said, even if she is being sincere. So, perhaps the author does not really believe every proposition in her book; again, this is a hurdle for representing the preface case as involving the author in a genuine inconsistency. Similarly,

[5] The preface paradox in its original form was formulated by D.C. Makinson [1965].
[6] See, e.g., Evnine [1999] and Roush [2010].
[7] See, e.g., Stalnaker [1984: 93–4], Weatherson [2005: 429–31], and Leitgeb [2014a]. Leitgeb, for example, proposes interpreting authors of books as merely asserting that the claims are highly probable and that, taken as a set, they are mostly true.

one may suspect that our verdict in the case tacitly depends upon assuming (as may be quite probable) that one of the individual propositions in the book is individually irrational, and thus that the case does not constitute a genuine counterexample to deductive closure principles.[8]

One could try to fix up the preface case to deal with these problems. However, I think it is less distracting to abstract away from the particular case, and simply to note the features that will generate the problems for the Deductive Closure Constraint (and its weaker cousins), since they are general and do not require any stories about prefaces or books. One has a number of individual beliefs which are each in some sense highly likely but with a small chance of error. As one puts together more and more of these beliefs, the risk of error aggregates, until it becomes overwhelmingly likely that at least one of them is mistaken. Then, the thought goes, it will be irrational to believe some logical consequence of the large set, if that logical consequence is true only if the whole large set is true. The simplest example is a long conjunctive proposition formed out of the individual beliefs in the set; but other examples are possible, too.

3. A New Way of Generating the Problem

It is often thought that whether the preface paradox is a devastating challenge to the Deductive Closure Constraint will turn on how one is inclined to think about the relationship between outright (binary) belief and credence (graded belief). The standard story is that, if one thinks of outright belief as reducible to credence, one will naturally expect the Deductive Closure Constraint to fail. For once we start thinking probabilistically, as we do when we focus on credence, we will start to worry about risk aggregation.[9] For example, suppose that one accepts the so-called 'Lockean thesis', according to which there is some threshold T such that, for any proposition p, a subject S believes p iff S's credence for p $\geq$ T. Then, in preface-type cases, one will have to violate the Deductive Closure Constraint if one is to remain probabilistically coherent.

Meanwhile, the story goes, if one thinks of outright belief as fundamental, and as irreducible to credence, one will be attracted to the Deductive Closure Constraint. If no arbitrarily high credence guarantees outright belief, it is thought, there is no obvious tension between probabilistic coherence of one's credences, on one hand, and deductive closure of one's outright beliefs, on the other.[10] Moreover, when one thinks in terms of outright beliefs, belief in a conjunction may seem to add no new commitment that belief in its individual conjuncts doesn't bring already. This can make it hard to see how it is that taking on the further conjunctive belief could be irrational if the existing beliefs involve no irrationality in and of themselves.

I will now show that one can use a preface paradox-type case to generate the problem for the Deductive Closure Constraint without making any assumptions about the reducibility of belief to credence. Instead, all we need are some very weak normative

[8] See, e.g., Ryan [1991].

[9] Thus Christensen [2004] and Sturgeon [2008], both deniers of the Deductive Closure Constraint, suggest that the attractiveness of such constraints depends to a great deal on how one thinks about the relationship between belief and credence. On the other side of the debate, defenders of the constraint, such as Douven and Williamson [2006], and Leitgeb [2014b], likewise present the challenge for closure as arising from a threshold model of belief.

[10] See, e.g., Stalnaker [1984] and Weatherson [2005].

assumptions about *rational* relationships between belief and credence. What we need specifically are the following claims:

> **Weak Fallibilism about Belief.** It's possible that
> (a) there are many (suitably independent) propositions such that, for each proposition p in this set, (i) S does not have credence 1 in p, and (ii) S (outright) believes p;
> (b) S is not radically probabilistically incoherent; and
> (c) in satisfying (a) and (b), S violates no requirement of rationality.

> **Belief-Credence Coherence.** Necessarily, rationality requires of S that, if S has a credence below 0.5 in p, then S does not (outright) believe p.

The first principle is labelled as a kind of *fallibilism* about belief because it affirms the rational permissibility of believing a non-trivial number of non-trivially independent propositions, without investing absolute certainty in these propositions (and without radical probabilistic incoherence). Henceforth, I'll just call it 'Weak Fallibilism'.

Careful readers will notice that Weak Fallibilism is somewhat vague, in its use of the words 'many', 'suitably', and 'high'. What is required, more precisely, is that there be enough propositions, that they be independent enough, and that S's credences for them be such that, in order to have a credence above 0.5 in their conjunction, S would have to be radically probabilistically incoherent.[11] That may sound somewhat demanding, but it is not. For example, if one has credence 0.9 in each of just 7 probabilistically independent propositions, the probabilistically coherent credence for their conjunction is 0.48; if one has credence 0.9 in each of 25 probabilistically independent propositions, the probabilistically coherent credence for their conjunction is 0.07. If one has credence 0.99 in each of 70 probabilistically independent propositions, the probabilistically coherent credence for their conjunction is 0.49; if one has credence 0.99 in each of 250 probabilistically independent propositions, the probabilistically coherence credence for their conjunction is 0.08.

Belief-Credence Coherence is, as the name suggests, a coherence requirement. It effectively imposes a minimum lower bound, 0.5, such that it is irrational to simultaneously believe a proposition and yet also have a credence for that proposition below the lower bound. However, it is crucial to see that this is not a reduction of belief to credence. The idea is neither that credence higher than 0.5 automatically counts as outright belief, nor that it is always rationally permissible to believe something for which one has a credence (or even a rational credence) higher than 0.5. Rather, the idea is just that it is irrational—because incoherent—to believe something for which one has a credence lower than 0.5. Moreover, Belief-Credence Coherence is 'wide-scope'.[12] That is to say, it does not give any special authority to your credences, saying that your credence determines what it is rational for you to believe. It does not say that, in some

[11] I cannot put an exact characterization on 'radically', which is also vague. The idea here is that the probabilistic incoherence of a set of beliefs admits of degrees. For example, if one assigns p a probability of 0.9 and q a probability of 0.9, and p and q are independent, it would be only slightly probabilistically incoherent to assign (p&q) a probability of, say, 0.8 (rather than 0.81), but it would be radically probabilistically incoherent to assign (p&q) a probability of, say, 0.08. (For an attempt to give a more precise characterization of this degreed notion of incoherence, with explanations and critiques of previous competing approaches, see Staffel [2015].) While the threshold for 'radically' is clearly vague, any reasonable precisification of 'radically' will do for the purposes of my argument. Note also that anyone who thinks that perfect probabilistic coherence is permissible will immediately agree, *a fortiori*, that some non-radically-probabilistically-incoherent states are permissible, irrespective of how we construe 'radically'.

[12] See, e.g., Broome [2007] on wide-scope rational requirements.

situation where you violate it, your outright belief rather than your credence is what should be revised. It just says that, one way or another, you should line up the two such that you do not combine a belief in p with a credence lower than 0.5 in p.

Now I'll spell out why Weak Fallibilism and Belief-Credence Coherence require rejection of the Deductive Closure Constraint with which we began. Suppose that we have a doxastic subject, Gina. Suppose that Gina (outright) believes a large number of propositions, and assigns each one a very high, but non-zero, credence. Suppose further that Gina is not radically probabilistically incoherent. Given suitable independence between the propositions she believes, Gina will, if she is not radically probabilistically incoherent, assign the conjunction of these propositions—call it c—a credence below 0.5. According to Weak Fallibilism, given these facts so far, it's possible that Gina violates no requirement of rationality. But if both Belief-Credence Coherence and the Deductive Closure Constraint were true, the facts specified so far would guarantee that Gina violates a requirement of rationality. For either Gina believes c or she doesn't. If Gina does believe c, she violates Belief-Credence Coherence, since she has a credence below 0.5 in c. If Gina doesn't believe c, she violates the Deductive Closure Constraint, since her beliefs entail c. So, by both of these cases, Gina violates a requirement of rationality. That contradicts what Weak Fallibilism said about Gina. So, Weak Fallibilism and Belief-Credence Coherence require the rejection of the Deductive Closure Constraint.

Neither Weak Fallibilism nor Belief-Credence Coherence involves any commitment to reducing outright belief to credence. So, we can get a serious problem for the Deductive Closure Constraint going without assuming anything about the reducibility of outright belief to credence. The challenge here carries over as well to the Deductive Closure Permission. According to Weak Fallibilism, Gina's existing body of beliefs is not irrational. However, even though her existing body of beliefs entails c, believing c in combination with her existing body of beliefs would (if she is not radically probabilistically incoherent) put her in violation of Belief-Credence Coherence and thus make her irrational. Thus, Gina is a counterexample to the Deductive Closure Permission.

In the next two sections, I will defend Weak Fallibilism and Belief-Credence Coherence, in turn.

4. Defending Weak Fallibilism

Weak Fallibilism is a bare possibility claim. It reflects the *possible permissibility* (not even the requirement) of both (a) displaying an extremely minimal humility about many of your beliefs, by not assigning them credence 1, and (b) not being radically probabilistically incoherent.

One way to deny Weak Fallibilism would be to hold that one may not permissibly have any non-trivially large number of non-trivially independent beliefs. I will simply assume that this is not an acceptable option. I do not know of anyone who has defended it.[13]

[13] One might argue that, since we tend to rely on just a few sources (memory, testimony, perception) for many of our beliefs, our beliefs should not be independent of each other to any large degree. (Thanks to Ram Neta here.) But although perfect probabilistic independence is an unrealistic simplifying assumption, propositions can still be fairly independent while being based on the same coarsely individuated sources. For many error-possibilities for individual propositions involve only local error, not the general failure of memory (or testimony, or perception) in all cases.

Another way to deny Weak Fallibilism would be to hold that rationality requires one to be radically probabilistically incoherent. Again, I don't know of anyone who has defended this, and I will assume that it is not an option. Some have argued that rationality does not require precise probabilistic coherence, or even that rationality requires one not to be precisely probabilistically coherent (since this would involve maximal confidence in logical truths).[14] But neither of those claims gets us to the claim that rationality requires radical probabilistic incoherence.

A somewhat more serious proposal for blocking Weak Fallibilism comes from the view that one cannot permissibly have outright beliefs in propositions in which one has non-1 credence. One way to motivate this view is to claim that belief just *is* credence 1.[15] If that is right, then one cannot have beliefs that one does not invest with credence 1, and so *a fortiori* cannot permissibly have such beliefs. So, Weak Fallibilism would be false. Interestingly, this suggests that the single view about the relationship between belief and credence that clearly evades the problem for deductive closure (as I have set it up) is itself a reductive view, albeit a very special kind of reductive view.[16] This is just the opposite of what is standardly assumed (cf. section 3).

Prima facie, the idea that belief is credence 1 looks unpalatable, for simple reasons. We believe a lot of things, but it seems that we are maximally certain of very few things. We report ourselves as believing, and we assert outright, a lot of things of which we would not say that there is literally *no* chance that they are false. Moreover, it's easy to think of situations that reveal our lack of maximal certainty in the things we believe—for example, bets at very bad odds.

However, Roger Clarke [2013] has recently defended a version of the view that belief is credence 1 which purports to get around these objections. His strategy involves appealing to what he calls 'sensitivism' about degrees of belief, according to which a subject's degree of belief changes 'from context to context, depending on the space of alternative possibilities' [ibid: 1]. It is important to see that although Clarke talks of 'context', he does not (I think)[17] mean to endorse any *semantic* thesis about the word 'belief' (or 'credence') and its semantic value changing as the conversational context of

[14] See, e.g., Christensen [2007] for discussion.

[15] One could hold the normative thesis that one is *rationally required* only to believe things to which one assigns credence 1, without embracing this descriptive reduction of belief to credence 1. However, it is the latter view that has been explicitly motivated in recent literature, and the former view on its own looks somewhat unmotivated. So, I focus on the latter view. However, the former view would be subject to strongly analogous objections to those that I give for the latter. Thanks to an anonymous referee here.

[16] Douven and Uffink [2003] argue that a similar solution can to the problem can be mimicked by a view on which one is to some degree *uncertain* about one's credences, as long as credence 1 is within the range of 'measurement error' from one's estimate of one's own credences. When one's estimates of one's credences in one's beliefs result in the probability of the conjunction of these propositions falling below a rationally acceptable threshold, one should then 'reinterpret' upwards one's (estimates of one's) credences—out of charity to oneself—so that the new probability of the conjunction exceeds this threshold. There are myriad questions one could ask about this view. Here it suffices to note that, as the set of one's independent beliefs reaches any non-trivial size, the credence level that this view requires us to 'interpret' ourselves as having in individual propositions will approach 1, getting very close to 1 very quickly. It also requires us to assume, without evidence, that we have systematically underestimated the values of our own credences, so as to affect the required 'reinterpretation' of our credences.

[17] Clarke writes elsewhere, in describing his view, that 'belief is context-sensitive … an agent may count as believing that p in one context but not another, without any change in her doxastic state between the two contexts' [forthcoming: 9]. This makes the thesis sound more semantic. But, on a semantic view, I do not understand how Clarke could say that whether a subject believes p depends on whether *the subject* is taking not-p possibilities seriously. And without such claims, Clarke's defence of his central claim that belief is credence 1 would be blocked.

the belief-*attributor* changes. Rather, the idea is that the subject's beliefs actually *change* depending on the salient 'space of alternative possibilities' for *her* qua subject. This space of alternative possibilities may be affected by factors that are, by the subject's own lights, evidentially irrelevant, such as changes in what is practically at stake, or in what sort of reasoning the subject is engaged in.

This idea of the sensitivity of doxastic states in general is not new with Clarke; it has been defended before, concerning outright belief.[18] What is perhaps new with Clarke is the idea that one's *credence* is sensitive to one's situation in a similar way (although Stanley [2005] hints, at times, at such a thesis about rational credence). On the picture for which others have argued previously, one's underlying credence, or level of confidence, stays constant across situations (barring apparent changes of evidence), and what changes is just whether that level of confidence suffices (either descriptively or normatively) for full belief given the situation.[19] Clarke's view is more radical. On his view, one has credence 1—and hence believes—a proposition when one does not take seriously any alternatives in some particular situation. But whether one takes alternatives seriously depends on the stakes, what is salient, and the kind of reasoning in which one is engaged. So, as one's situation changes, one shifts from assigning a proposition credence 1 (and hence believing it) to assigning it credence < 1 (and hence not believing it). Treating not only belief, but also credence, as sensitive allows Clarke to maintain that belief just is credence 1.

I offer four objections. First, and most simply, our ordinary belief-attributing practices allow us to self-attribute beliefs even while explicitly acknowledging a possibility of error *simultaneously and in the same context*. For example, it sounds completely fine to say something of the form 'I believe that P, but I'm not sure' or 'I believe that P, but I might be wrong' or 'I believe that P, but it might be that Q,' where the proposition that Q is an incompatible alternative to the proposition that P.[20]

Second, consider some case where the stakes change and I start taking some counter-possibility seriously. Take some ordinary proposition I believe, like the proposition that I was born in the month of August. Suppose now that I'm offered some very unfavourable bet on this proposition—I win $1 if it is true and lose everything I own if it is false. On Clarke's account, I began with credence 1, since I do not ordinarily take seriously the possibility that I was not born in the month of August—but when I am offered the bet I start taking those possibilities seriously, my credence changes, and I lose my belief. Now, as I understand it, a credence is (or at least commits one to) a kind

[18] See Weatherson [2005], Ganson [2008], and Leitgeb [2014b]; for a critique and articulation of a compelling alternative, see Ross and Schroeder [2014]. Weatherson and Leitgeb seem to think that this sensitivist view can help us in defending deductive closure against preface paradox-type worries (see also Hawthorne [2003: 183]). But it is not clear how the sensitivist view on its own helps in resisting my specific argument against the Deductive Closure Constraint, unless this is to be achieved by embracing the result that one may not permissibly, at any single point in time, have any non-trivial number of non-trivially independent beliefs. (This appears to be a consequence of Leitgeb's view as it stands, although he does not make this explicit.)

[19] Ganson [2008: esp. 455] is especially clear on this point.

[20] It might be replied that sentences of the form 'I believe that P' are not really belief-reports, but are rather 'hedged assertions' that P (cf. Murray [2014]). Perhaps the way to express a belief that P is simply to assert that P; and assertions of the form 'P, but it might be that Q' (where Q is incompatible with P) do not sound good. I am sceptical that sentences of the form 'I believe that P' are never belief-reports. Moreover, there are ways of making it clear that one is reporting a belief, for which the conjunctions still sound good: for example, one can say: 'My own view is that P, but it might be that Q' or even 'I firmly believe that P, but it might be that Q.' Finally, there are alternative explanations of the badness of 'P, but it might be that Q' that do not appeal to some incompatibility between belief and recognizing a chance of error (see Worsnip [forthcoming-a]).

of estimate of the probability or likelihood of a proposition. Consequently, when one changes one's credence for some proposition, one is thereby committed to thinking that *either* the probability of the proposition has changed *or* one made a mistake in one's previous estimate of that probability. But, in the August case, neither option seems plausible. It is very unnatural for me to think that, after I have been offered the bet, the probability that I was born in August has changed. But nor does it seem that, after I am offered the bet, I need to think that I made a *mistake* in my original assessment of the probability that I was born in August.[21] The much more natural thing to say, after I am offered the bet, is that I take the probability that I was born in August to be the same as I always did take it to be. What has changed is only whether that probability suffices for reliance on the proposition. This is just the standard decision-theoretic picture: credences remain stable across practical situations, providing a baseline from which to calculate the expected utility of different courses of action as the costs change.

Third, recall that, on Clarke's view, if I am treating something as certain and ignoring the counter-possibilities, I count as having credence 1 in the circumstances. However, in order to even work out whether I may treat a proposition as certain and ignore counter-possibilities, I need some baseline estimate of the probability of that proposition. I know that, ordinarily, I can just treat as certain the proposition that I was born in August; I also know that in special circumstances, like that of the unfavourable bet, I cannot do so. But in order to be in a position to know either of those things, I need a baseline estimate of the probability that I was born in August. This baseline estimated probability *explains* why I may treat the proposition as certain in ordinary circumstances, as well as explaining why I may not treat the proposition as certain in special circumstances. So, we must admit that, in addition to talking of *treating* a proposition as certain, and ignoring counter-possibilities, we also need what is, in one sense, a more fundamental notion of my actual estimate of the probability of the proposition.[22] But that latter notion, I think, just is the one that decision theorists and epistemologists mean to pick out in talking of 'credence'. Even if one reserves the term 'credence' for the former notion, one could still substitute the latter notion for 'credence' in our argument against deductive constraints. So, Clarke's view does not succeed in evading that argument.

Fourth, I find it hard to understand how Clarke's view applies to non-occurrent beliefs. Take the classical preface case. Clarke [forthcoming] says that when the author is at work on the body of the book she believes the claims in the book—but that when she is writing the preface she attends to the possibility of error, and thus does not believe the claims in the book (hence avoiding synchronic inconsistency with her claim that she has made errors). But what about times at which she is not actively at work on

[21] Of course, this second diagnosis may be right in *some* cases where a change in stakes causes me to revisit my original assessment, to attend to some possibilities that I had been overlooking, and to adjust an unreasonably overconfident credence (cf. Nagel [2008]). But this had better not be Clarke's general story about what happens for ordinary beliefs. For if changes in the stakes cause us to realize that our assignments of credence 1 were *unreasonable*, then on Clarke's account, by the same token, the beliefs themselves are unreasonable (since belief just is credence 1). So, generalized, this view would commit Clarke to saying that all of our ordinary beliefs are unreasonable.

[22] This may map onto what Wedgwood [2012] calls 'practical credence', as opposed to 'theoretical credence'. He identifies full belief with being disposed to have *practical* credence 1 in normal circumstances. Unlike Clarke's view, this does not block our argument against deductive closure. We can just use theoretical, rather than practical, credence in the argument.

her book at all—attending to the claims neither in the book nor in the preface? I take it that, at any particular point in time, we have lots of beliefs to which we are not attending. Yet it is hard to see what Clarke should say about such non-occurrent beliefs. It seems arbitrary to say either that at such times the author does not believe the claims in her book, or that at such times she does not believe the claim that her book contains errors.

I can think of one alternative way of trying to resist Weak Fallibilism, which is to claim that preface paradox cases are, in effect, rational dilemmas. On this view, although one rationally ought to have lots of beliefs to which one assigns a non-1 credence, and although one rationally ought not to be radically probabilistically incoherent, one also rationally ought to obey the Deductive Closure Constraint. Unfortunately, that means that one cannot be rational.

I think that this position is best avoided if at all possible. As I pointed out earlier, preface paradox situations are not limited to special cases involving books and prefaces; they can arise for any large set of beliefs. So, this approach does not posit rational dilemmas only in some unusual cases; it posits them all over the place. Moreover, allowing rational dilemmas makes it much harder to make progress in a theory of rationality generally. One important way of arguing against particular putative requirements of rationality is to show that one has to violate them if one is to obey more compelling, non-negotiable requirements. But this method of argument assumes that there are not rational dilemmas. If we are willing to allow rational dilemmas to spring up all over the place, this method will not work.

That is not to deny that in some sense there are conflicts between normative ideals that would, in some sense, be *good to fulfil*.[23] No doubt, in one sense life would be easier if the Deductive Closure Permission held, and we could always infer the logical conclusions of our existing beliefs without any need to worry. No doubt, in one sense things would be neater if our beliefs were deductively closed. But this does not mean that we are irrational if, noticing that we are not in an ideal world, we knowingly violate the Deductive Closure Constraint. The view that there are rational dilemmas *proper* is the strong claim that, no matter what one does in such a situation, one is *irrational*: that there is no rational way to resolve such a situation. I do not think that this strong judgment is warranted in preface-type cases. Having beliefs that are not deductively closed may in one way be regrettable, but that does not make it irrational.

5. Defending Belief-Credence Coherence

One might wonder why Belief-Credence Coherence mentions credence 0.5 in particular. The answer is that, at the point where one has a credence lower than 0.5 in a proposition, one judges the proposition more likely to be false than true. This attitude, I claim, is rationally incompatible with belief. Part of what it is to believe something is to take it to be true—where the relevant contrast is with its being false. It is incoherent both to take something to be true and to judge it to be more likely false than true.[24]

[23] Christensen [2007], who is often read as arguing for rational dilemmas in a strong sense, sometimes sounds like he is only committing himself to this weaker claim.

[24] I actually also think that it is irrational to believe something and to have a credence of *exactly* 0.5 in it, since the latter attitude still involves taking it to be no more likely to be true than false. But I need not rely on this.

The incoherence involved in violating Belief-Credence Coherence can be brought out, as some philosophers have noted [Kaplan 1996: 142–3; Christensen 2004: 48–9], by considering the Moore-paradoxical flavour of utterances that report the two mental states jointly. It seems incoherent to say 'P, but it's more likely that not-P' or to say 'I believe that P, but it's more likely that not-P.' However, Mark Kaplan [1996] has defended a view on which one can (rationally) believe p even when one has a credence lower than 0.5 in p. He offers an interesting argument by analogy for there being no incoherence in such a combination of mental states [ibid: 144–5]. The analogy is a choice where one is offered either a certain $5, or a bet whereby one wins $10 if some proposition p is true but loses $1000 if p is false. Suppose that one is 0.95 confident in p. Then the expected value of the bet is lower than the certain $5. In that case, one might pick the certain $5. Nevertheless, one is 0.95 confident that taking the bet would maximize one's actual winnings.

Kaplan points out, rightly, that there is no incoherence here. He then claims that, in preface-like cases, combining a belief in p (where p is, for example, a long conjunction of propositions) with high confidence that p is false is just the same. This, I do not find convincing. To believe p is to believe that p is true. On the other hand, to choose an option is clearly not to believe that this option will have the best actual outcome of all the available options. Precisely because one makes decisions under risk, one can—to protect oneself from a chance of disaster—hedge one's bets and choose something that one thinks probably won't end up having the best outcome of the available options. That is exactly what happens in Kaplan's betting case. Since choosing the option does not necessarily involve thinking that it will have the best possible outcome, it does not clash with thinking that some other incompatible option is more likely to have the best possible outcome. But since believing something does involve believing it to be true, it does clash with thinking that some incompatible proposition (such as its negation) is more likely to be true.

Relatedly, in the case of the bet, even though one cannot say that declining the bet is more likely to have the best actual outcome of the available options, one can say that declining the bet has the best *expected* outcome of the available options (in the sense of 'expected' used in 'expected utility' theory).[25] There is nothing analogous that one can say in the case of the belief that one judges to be more likely false than true. It's not as though one can say that, even though it is more likely to turn out false than true, the 'expected truth-value' of the proposition is still *true*. Nor is belief in p somehow a way of 'hedging one's bets', as it is in the betting case. If anything counts as a 'hedge' in cases where one judges p very likely to be false, it is suspending judgment about p—not *believing* p.

Indeed, Kaplan's own view of belief does not, on the face of it, seem hospitable to belief in things that one judges to be more likely false than true. On Kaplan's view, which he labels the 'assertion view', 'you count as believing p just if, were your sole aim to assert the truth (as it pertains to p), and your only options were to assert that p, assert that not-p, or make neither assertion, you would prefer to assert that p' [ibid: 109]. One would think that, in any case whereby one judges that p is more likely to be

[25] As simple examples like Kaplan's show, taking an option to have greater expected utility than another option does not necessarily involve 'expecting' that it will have a better actual outcome. The former estimate takes into account not just the relative probabilities of the possible outcomes but also their *degrees* of goodness or badness.

false than true, to the extent that one's sole aim is to assert the truth as it pertains to p, one would (if one is rational) prefer *not* to assert p!

Kaplan apparently thinks that this is mitigated by the fact that one is aiming to tell a comprehensive story about the world [ibid.: 111]. The thought seems to be this: there are lots of propositions, each of which I would prefer to assert rather than to assert its negation or to make neither assertion. So, given that I aim to tell a comprehensive story about the world, plus the assertion view of belief, I will count as believing many of these things. The problem, however, comes in getting from this claim to the claim that I believe the logical consequences of these individual propositions (for example, their conjunction). It seems to me that Kaplan must be moving from the claim that I would be willing to assert each proposition individually to the claim that I would be willing to assert their conjunction. But that just doesn't follow.

Indeed, Kaplan's picture of belief *actively reinforces* the idea that there is a big, important, difference between believing a set of individual propositions and believing their conjunction. On some views of belief, it may be hard to put one's finger on what it is to believe a set of individual propositions without believing their conjunction. But it is *easy* to put one's finger on what it is to be disposed to assert any one of a set of individual propositions without being disposed to assert their conjunction. And it is likewise easy to see why one might rationally be in that position, if one's sole aim is to assert the truth as it pertains to each proposition in question. For, as preface paradox cases teach us, it can be that, with respect to each individual proposition, I am very likely to be asserting the truth but, with respect to the conjunctive proposition, I am not. So, I think that the assertion account of belief actually *strengthens* the case against deductive closure.

6. Comparisons with Other Ways of Generating the Problem

Before I close, I will compare my strategy with two other ways of generating the problem that, like my approach and unlike many standard approaches, do *not* assume the reducibility of belief to credence. In both cases, I will suggest, my method of generating the problem uses more minimal assumptions, closing off more possible lines of resistance.

First, consider the approach of James Hawthorne and Luc Bovens [1999]. They reject the Lockean thesis as a descriptive reduction, but they accept a normative variant of it according to which an *ideally rational* subject will have some threshold T such that, for any proposition p, she believes p iff her credence for p $\geq$ T. They avoid reductivism, since they nevertheless allow that many subjects do not satisfy this requirement. Their paper shows neatly that, if one accepts the normative variant of the Lockean thesis, one can generate rational failures of deductive closure.

However, in order for this demonstration to work, it's crucial that they assume that an ideal subject has a *single* threshold for believing, one that is fixed regardless of context or subject-matter. That claim is, at a minimum, very controversial; indeed, I think that it is false. It is too demanding to say that subjects must have the same belief-threshold for all propositions regardless of subject-matter. Consider lottery cases. It seems that the credence-threshold we set for 'My ticket will lose the lottery' is *much* higher than it is for most ordinary propositions; and this is not obviously irrational. Indeed, Hawthorne and Bovens themselves express sympathy for a pragmatic view of belief on which 'belief might well be just a matter of [one] stipulating a confidence level that [one] finds high enough to be of special significance' [ibid: 256]. But surely what is of special significance might vary depending on context and subject-matter.

My approach avoids the questionable assumption here. While I impose a requirement that a rational subject *not* believe something for which her credence < 0.5, I do not impose any threshold such that a rational subject must *believe* a proposition to which she assigns a credence of at least that threshold. Nor do I require that a subject have any such fixed personal threshold. So, I allow that a rational subject could have a variable threshold for outright belief depending on context or subject-matter—as long as it never dips below 0.5—or indeed no specific threshold at all which suffices for her believing, even in particular contexts.

Second, one might imagine an approach which does everything, not in terms of credence, but in terms of evidential probability. The claim would be that, since one's evidence does not generally assign an evidential probability of 1 to the individual propositions one believes, and since evidential probability obeys probabilistic constraints, one's evidence assigns a very low evidential probability to long conjunctions of these beliefs. Then, on the assumption that rational requirements cannot come into conflict with believing what the evidence supports, one can plausibly generate failures of the Deductive Closure Constraint.[26]

Unfortunately, it is not clear that this assumption is correct. The Deductive Closure Constraint is a putative coherence requirement of rationality, in the sense that it specifies structural relations that are (putatively) required to hold between one's different mental states for a subject to be fully rational. On some views, there are cases where such coherence requirements do come into conflict with responding to one's evidence.[27] A defender of the Deductive Closure Constraint might allege that preface paradox-type cases are also cases of this sort. On this view, showing that fulfilling the Deductive Closure Constraint can lead to a failure to believe what one's evidence supports does not suffice for showing that the Deductive Closure Constraint is false. For such cases may be instances of conflict between coherence requirements and evidence-responsiveness. By contrast, my challenge to the Deductive Closure Constraint comes from *within* the sphere of coherence requirements of rationality—Belief-Credence Coherence is a coherence requirement in the same sense that the Deductive Closure Constraint purports to be. No assumption about the impossibility of conflict between coherence requirements and evidence-responsiveness is required for my strategy.

Furthermore, the assumption that one's evidence does not generally assign an evidential probability of 1 to the individual propositions that one believes is disputed by Timothy Williamson [2000], who claims that your knowledge has epistemic probability 1 for you.[28] On this view, if you really know every one of a large number of propositions, there is no aggregation of risk in terms of evidential probability when it comes to their conjunction.[29] If one combines this with a bold interpretation of Williamson's 'knowledge-first' epistemology, on which a rationally ideal subject will believe only that

[26] So, Easwaran and Fitelson [2015: 64–73] argue that deductive closure and consistency requirements conflict with the norm enjoining one to believe what one's evidence supports, and so must be rejected.

[27] I defend this view elsewhere [forthcoming-b]. See also Christensen [2007, 2013] and Fogal [manuscript]. Fogal explicitly suggests that preface paradox cases may be instances of this kind of conflict.

[28] See Williamson [2009] for application of this idea to preface-style cases.

[29] Williamson [2009] tries to explain away the counterintuitive sound of this by pointing out that, even though you know any conjunction built out of conjuncts that you know (and such a conjunction has probability 1), you may fail to know that you know many of the conjuncts, in which case the evidential probability *that you know* the conjunction will be very low, which is why it seems to you that you don't know the conjunction. See Ryan [1991: 297] and Smith [2010] for other views of evidential probability that are friendly to deductive closure constraints.

which she is in a position to know, this view seems to withstand a challenge to deductive closure framed in terms of evidential probability. Again, my approach avoids this assumption. Because I do not employ the notion of evidential probability, I do not need to make any assumptions about it. Instead, I employ the notion of credence (without assuming the reducibility of belief to credence). Admittedly, someone who holds that knowledge has epistemic probability 1 might also hold that belief is credence 1, but that would return us to the problems with the view of belief as credence 1 that were already explored in section 4.[30]

In sum, while it may have looked like the objection to deductive closure (both the constraint and the permission) from preface paradox-style cases rests upon a particular view of belief, credence, or evidence, the version of the preface paradox developed here shows that this is a mistake. As such, things look worse for deductive closure than its defenders have supposed.[31]

References

Broome, John 2007. Wide or Narrow Scope? *Mind* 116/462: 359–70.
Christensen, David 2004. *Putting Logic in its Place: Formal Constraints on Rational Belief*, Oxford: Oxford University Press.
Christensen, David 2007. Does Murphy's Law Apply in Epistemology? Self-Doubt and Rational Ideals, in *Oxford Studies in Epistemology, vol. 2*, ed. Tamar Szabó Gendler and John Hawthorne, Oxford: Oxford University Press: 3–31.
Christensen, David 2013. Epistemic Modesty Defended, in *The Epistemology of Disagreement: New Essays*, ed. David Christensen and Jennifer Lackey, Oxford: Oxford University Press: 77–97.
Clarke, Roger 2013. Belief is Credence One (in Context), *Philosophers' Imprint* 13/11: 1–18.
Clarke, Roger forthcoming. Preface Writers are Consistent, *Pacific Philosophical Quarterly*.
Douven, Igor and Jos Uffink 2003. The Preface Paradox Revisited, *Erkenntnis* 59/3: 389–420.
Douven, Igor and Timothy Williamson 2006. Generalizing The Lottery Paradox, *The British Journal for the Philosophy of Science* 57/4: 755–79.
Easwaran, Kenny and Branden Fitelson 2015. Accuracy, Coherence, and Evidence, in *Oxford Studies in Epistemology, vol. 5*, ed. Tamar Szabó Gendler and John Hawthorne, Oxford: Oxford University Press: 61–96.
Evnine, Simon J. 1999. Believing Conjunctions, *Synthese* 118/2: 201–27.
Fogal, Daniel manuscript. Rationality and Requirements.
Ganson, Dorit 2008. Evidentialism and Pragmatic Constraints on Outright Belief, *Philosophical Studies* 139/3: 441–58.
Harman, Gilbert 1986. *Change in View: Principles of Reasoning*, Cambridge, MA: The MIT Press.
Hawthorne, James and Luc Bovens 1999. The Preface, the Lottery, and the Logic of Belief, *Mind* 108/430: 241–64.
Hawthorne, John 2003. *Knowledge and Lotteries*, Oxford: Clarendon Press.
Kaplan, Mark 1996. *Decision Theory as Philosophy*, Cambridge: Cambridge University Press.
Leitgeb, Hannes 2014a. A Way Out of the Preface Paradox? *Analysis* 74/1: 11–15.
Leitgeb, Hannes 2014b. The Stability Theory of Belief, *The Philosophical Review* 123/2: 131–71.
Makinson, D.C. 1965. The Paradox of the Preface, *Analysis* 25/6: 205–7.

[30] I do not think that this latter view need *automatically* come along with the former. At least in his early work, Williamson seems hesitant to identify belief with credence 1 [2000: 213, 2005: 685], although these passages are open to multiple interpretations.

[31] For helpful discussion of these issues, I'm grateful to an audience at the Joint Session of the Aristotelian Society and the Mind Association, and more particularly to Luc Bovens, Julianne Chung, Stephen Darwall, Keith DeRose, Julien Dutant, Tamar Gendler, Jeremy Goodman, Daniel Greco, Mark Maxwell, Ram Neta, Aaron Norby, Bernhard Salow, and especially Bruno Whittle. I am also grateful for comments from two anonymous referees for *AJP* and from the editor thereof.

Nagel, Jennifer 2008. Knowledge Ascriptions and the Psychological Consequences of Changing Stakes, *Australasian Journal of Philosophy* 86/2: 279–94.
Ross, Jacob and Mark Schroeder 2014. Belief, Credence, and Pragmatic Encroachment, *Philosophy and Phenomenological Research* 88/2: 259–88.
Roush, Sherrilyn 2010. Optimism about the Pessimistic Induction, in *New Waves in Philosophy of Science*, ed. P.D. Magnus and Jacob Busch, Basingstoke, UK: Palgrave MacMillan: 29–58.
Ryan, Sharon 1991. The Preface Paradox, *Philosophical Studies* 64/3: 293–307.
Smith, Martin 2010. What Else Justification Could Be, *Noûs* 44/1: 10–31.
Staffel, Julia 2015. Measuring the Overall Incoherence of Credence Functions, *Synthese* 192/5: 1467–93.
Stalnaker, Robert 1984. *Inquiry*, Cambridge, MA: The MIT Press.
Stanley, Jason 2005. *Knowledge and Practical Interests*. Oxford: Oxford University Press.
Sturgeon, Scott 2008. Reason and the Grain of Belief, *Noûs* 42/1: 139–65.
Weatherson, Brian 2005. Can We Do Without Pragmatic Encroachment? *Philosophical Perspectives* 19/1: 417–43.
Weatherson, Brian 2012. Knowledge, Bets, and Interests, in *Knowledge Ascriptions*, ed. Jessica Brown and Mikkel Gerken, Oxford: Oxford University Press: 75–103.
Wedgwood, Ralph 2012. Outright Belief, *Dialectica* 66/3: 309–29.
Williamson, Timothy 2000. *Knowledge and Its Limits*, Oxford: Clarendon Press.
Williamson, Timothy 2005. Knowledge and Scepticism, in *The Oxford Handbook of Contemporary Philosophy*, ed. Frank Jackson and Michael Smith, Oxford: Oxford University Press: 681–700.
Williamson, Timothy 2009. Replies to Critics, in *Williamson on Knowledge*, ed. Patrick Greenough and Duncan Pritchard, Oxford: Oxford University Press: 279–384.
Worsnip, Alex forthcoming-a. Possibly False Knowledge, *The Journal of Philosophy*.
Worsnip, Alex forthcoming-b. The Conflict of Evidence and Coherence, *Philosophy and Phenomenological Research*.