

*Why Have Consistent and Closed Beliefs,
or, for that Matter, Probabilistically Coherent Credences?**

What should we believe? At very least, we may think, what is logically consistent with what else we believe, and what logically follows from what else we believe. Or if the question is which degrees of belief we should have, degrees of belief that, at very least, conform to the probability axioms. As these answers reflect, we seem to accept “norms of pure coherence”: norms which demand no more and no less than that our beliefs, or degrees of belief, be formally—logically or probabilistically—coherent. Two examples are:

Simple Consistency: We are rationally required not both to believe p and to believe not p .
and

Simple Closure: If q is a logical consequence of p , we are rationally required if we believe p , to believe q .

Norms of pure coherence are, in a sense that John Broome’s pathbreaking work has made current, *wide-scope*: indeed, as I will put it, *strongly wide-scope*.¹ There is no specific belief that one must have, or that one must lack, and (unless the content in question is logically true or logically false) no specific degree of belief that one must have, or that one must lack, in order to satisfy a norm of pure coherence. By making suitable adjustments in one’s other attitudes, one can conform to the norms of pure coherence, with or without that particular belief or degree of belief. Norms of pure coherence are also, again in a sense that Broome has made current, “requirements of rationality.” Whether one satisfies a requirement of rationality, in having or lacking a particular belief, depends entirely on whether one has or lacks, or had or lacked, some other beliefs. We need no other information, but information about one’s beliefs, in order to

* [Thanks to audiences of drafts of parts of this paper at Amherst, UCL, Valencia.]

¹ [Refs. to Broome.]

know whether one is conforming to a norm of pure coherence. Finally, norms of coherence are independent of the evidence. Whatever the evidence is, they say, we ought to have logically consistent and closed beliefs, or probabilistically coherent degrees of belief. Whatever the evidence is, if our beliefs, or degrees of belief, are formally coherent in these ways, then we believe as we ought, along at least one dimension.

From the standpoint of a certain kind of philosophical reflection,² norms of pure coherence can seem the only unproblematic answers to the question, “What should we believe?” Substantive norms of evidence, by contrast, seem perennially contestable. We can say that *this* experience is evidence, of *this* strength, that *that* is true. But how is this anything more than ratification of our own intuitions? All that we can be sure of is that our beliefs ought to be logical (or as austere Bayesians will say) that our degrees of belief ought to be probabilistic. Yet if we consider the matter from the different, perhaps more naïve, standpoint of first-person deliberation about what to believe, it is the norms of pure coherence that seem puzzling. For when we ask, from the standpoint of first-person deliberation, “What to believe?” the answer is simple and exhaustive: what the evidence suggests is true and not what the evidence suggests is false. Or, if we ask which degrees of belief should we have: higher degrees in what the evidence suggests is true and lower degrees in what the evidence suggests is false.² What room is there, within this standpoint, for the different answers to the question “What to believe?” that the norms of pure coherence embody? The deliberative standpoint looks out on the world. It craves

² Of course, we can ask the related, but different questions, “On what questions should we try to arrive at beliefs, one way or the other?” and “How much evidence should we have, before arriving at beliefs, one way or the other, on these questions?” When we ask these questions, then other answers flood in: the instrumental or intrinsic value of the knowledge that we might achieve, and the instrumental or intrinsic disvalue of the error into which we might stray. Yet once it is settled which questions to settle, and once it is settled how much evidence will settle them, then those questions are to be settled simply on the basis of that evidence. From that point on, what matters is simply shunning falsity and cleaving to truth.

evidence; it seeks truth and shuns falsity. Why should it cast its gaze inward on itself? Why should it care about how its beliefs are arranged with one another?

One's initial answer may be that the norms of pure coherence are a way of following the evidence toward what is true and away from what is false. But—as the negative thread of this paper argues—this is not so. Norms of pure coherence are, as they appear to be, unmoored from the evidence. Seen in this light, norms of pure coherence seem to have something like the status of “pragmatic” considerations, such as the boost in confidence or divine favor that a belief might bring. Or, rather, something less than this status. For these pragmatic considerations are clearly important considerations, which are relevant to other questions, if not to the question what to believe. The norms of pure coherence, by contrast, seem scarcely important at all, like rules of mental housekeeping, or the expression of a fetish for a certain kind of psychic pattern.³

The positive thread of the paper then tries to explain what the negative thread makes mysterious: why we nevertheless find norms of pure coherence so natural. While logic does not directly govern our beliefs, as norms of pure coherence suggest, logic does structure the evidence, which does directly govern our beliefs. As a result, the evidence often—although importantly not always—supports patterns of belief that are logically consistent and closed. And so norms of pure coherence can seem plausible. Once we see this, I suggest, the paradoxes of the preface and the lottery, in which the evidence supports illogical patterns, should no longer seem so paradoxical. Nor should they seem grounds for abandoning a framework of full belief for one of graded degrees (whatever other, possibly compelling, reasons there might be for this change of focus), or for replacing norms of pure logical coherence for norms of pure probabilistic coherence. For norms of pure probabilistic coherence are problematic in much the

³ [Also a puzzle about why norms of pure coherence should apply to beliefs, and not, say, imaginings, if they are not animated by our interest in truth and falsity. Nishi Shah's point.]

same ways. The lesson, I suggest, is simply that there are no norms of pure coherence, logical or probabilistic. There is just the evidence and how best to respond to the evidence in order to believe, or believe to a higher degree, what is so and not to believe, or to believe to a lesser degree, what is not.

1. Simple consistency: the positive thread

It will be easiest to begin with the positive thread, applied to the simplest case. Why does the norm of pure coherence:

Simple Consistency As Such: One is rationally required not both to believe p and to believe not p.

seem plausible? In part, I think, because it appears to explain:

First Observation about Simple Consistency: If one believes that p and believes that not p, then one violates some norm.

[Broome on being entirely as you ought to be.] In this section, I will try to argue that the First Observation can be explained without appealing to Simple Consistency As Such.

When does epistemic reason permit you to believe that p? A rough answer is that it is when the evidence that p is sufficient. But this answer is at best incomplete. How much evidence is “sufficient” is, it seems, at least in part a comparative matter. After all, no matter how much evidence—supposing that there is some absolute measure of this—there is that p, if there is even stronger evidence that not p, then, it seems, intuitively, that epistemic reason does not permit one to believe that p. And if there is equally strong evidence that p and that not p, then it also seems, intuitively, that epistemic reason does not permit one to believe that p. If the evidence is equally balanced, then, the standard view is, one ought to suspend belief. In sum, intuitively, it seems that:

Stronger Evidence: Epistemic reason permits one to believe something only when the evidence that it is the case is stronger than the evidence that it is not the case.

Note that this is only a necessary condition for epistemic permission. It is compatible with there being other necessary conditions: e.g., that the evidence is stronger by a sufficient margin or that there is sufficient evidence in some absolute sense.

If Stronger Evidence is true, then epistemic reason can account for the First Fact. For from Stronger Evidence it follows that:

whenever epistemic reason permits one to believe that p, epistemic reason does not permit one to believe that not p, and vice-versa.

Suppose epistemic reason permits one to believe that p. Then, by Stronger Evidence, the evidence that p is stronger than the evidence that not p. Therefore, the evidence that not p cannot be stronger than the evidence that p. Therefore, by Stronger Evidence, epistemic reason does not permit one to believe that not p. And so it follows that:

In having simply inconsistent beliefs, in believing that p and believing that not p, one is invariably violating epistemic reason.

If epistemic reason does not permit one to believe that p, then one is violating epistemic reason by believing p. And if epistemic reason does permit one to believe that p, then it does not permit one to believe that not p, so one is violating epistemic reason by believing not p. The “norm” referred to in the First Observation is simply the relevant verdict of epistemic reason: the verdict that pertains to one’s actual situation.

Here’s another way of viewing the matter. In Stronger Evidence, we have one necessary condition for epistemic reason’s permitting the belief that p. Leaving open the possibility that

there are other necessary conditions, we can distinguish three jointly exhaustive evidential situations:

- (i) The evidence that p is stronger than the evidence that not p and the other necessary conditions for permitting a belief that p obtain.
- (ii) The evidence that not p is stronger than the evidence that p and the other necessary conditions for permitting a belief that not p obtain.
- (iii) None of the above.

In each evidential situation, epistemic reason delivers the following pair of verdicts, in light of (2):

- (i) Epistemic reason permits one to believe that p. Epistemic reason requires one not to believe that not p.
- (ii) Epistemic reason permits one to believe that not p. Epistemic reason requires one not to believe that p.
- (iii) Epistemic reason requires one not to believe that p. Epistemic reason requires one not to believe that not p.

Since these evidential situations are jointly exhaustive, one is always in one of these situations. So one is always subject to one of these pairs of verdicts. For each pair of verdicts, believing that p and believing that not p violates at least one verdict in the pair. Thus, in believing that p and believing that not p, one always violates some verdict of epistemic reason. *Which* verdict or verdicts one violates, of course, depends on which evidential situation one is in. Nevertheless, this is enough to explain the First Observation.

These verdicts of epistemic reason are narrow-scope norms; they require or forbid particular beliefs. And they are not pure coherence norms. Having a simply consistent pattern of

belief is necessary, but not sufficient for satisfying a given pair of these verdicts. One must also have the *specific* consistent pattern of belief that the verdicts call for. Nevertheless, it is very easy to confuse these requirements of epistemic reason with Simple Consistency As Such, and this may account for much of the latter's appeal. Since one of the three evidential conditions always obtains, it is always the case that:

- (a) Either (epistemic reason permits one to believe that p, but does not permit one to believe that not p),
or (epistemic reason permits one to believe that not p, but does not permit one to believe that p),
or (epistemic reason does not permit to believe that p, and does not permit one to believe that not p).

One might sum this up by saying:

- (b) Either (one is not permitted by something or other to believe that p)
or (one is not permitted by something or other to believe that not p).

If we then allow the "not permitted" to take scope over the whole disjunction, rather than over simply the individual disjuncts, then we get something that looks like Simple Consistency As Such:

- (c) One is not permitted by something or other (to believe p and to believe not p).

But (b) and (c) are quite different claims. While one can always satisfy the requirement (c) by believing that p, (b) might be true precisely because one is required not to believe that p.

Now, it might be suggested, first, that we may introduce disjunctions within the scope of deontic operators, and, second, that by doing so, we can derive (c) from (b). The alleged upshot is that we cannot appeal to Stronger Evidence to explain away Simple Consistency As Such,

because Stronger Evidence implies Simple Consistency As Such. Like others, I see little to recommend this rule of inference, which delivers such principles as, “One ought always either to do what one has most reason to do, or to seek the destruction of all that is good and holy.”⁴ But I need not resist the rule of inference too strenuously. My aim is to explain away the appeal of Simple Consistency As Such, an appeal which is not shared by “junk” norms, like “One ought always either to do what one has most reason to do, or to seek the destruction of all that is good and holy,” or “One ought always to believe what the evidence supports, or dance the Hokey-Pokey,” which have no basis but this sort of inference. If you accept this rule of inference, then simply add to “Simple Consistency As Such” the qualification, “understood as something other than a norm justified only by this rule of inference.”

A more interesting protest is that we have overlooked a role for Simple Consistency As Such that Stronger Evidence cannot explain. This role emerges in nonideal cases, in which we suppose that the agent does not comply with the verdicts of epistemic reason. Consider a case in which you presently believe that *p* and believe that not *p*. Although, let us suppose, epistemic reason requires you to believe that *p* and not to believe that not *p*, you instead drop the belief that *p* and continue to believe not *p*. Isn't it right to say that you have satisfied some norm? Isn't it something at least to have replaced your inconsistent beliefs with some consistent pattern? This norm cannot be one of epistemic reason. As far as epistemic reason is concerned, your position

⁴ Cite Ross, Raz. Of course, this principle does not imply that if you seek the destruction of all that is good and holy, then you do everything that you ought to do. For you still fail to satisfy the nondisjunctive requirement to do what you have most reason to do. And as we can explain pragmatically why we are unlikely to assert the obvious truth “Either you ought to do what you have most reason to do, or you ought to seek the destruction of all that is good and holy,” perhaps too we can explain pragmatically why we are unlikely to assert, even if it is true, “You ought either to do what you have most reason to do or to seek the destruction of all that is good and holy.” But these are scarcely reasons *for* accepting the rule of inference. The only reason for it, it seems, is that it simplifies the formal semantics for “ought.”

is worse than it was, for you have gone from having one belief that epistemic reason requires and having one that it prohibits, to lacking the belief that epistemic reason requires but still having the one that it prohibits. The norm in question seems to be one of rationality. There is a kind of irrationality involved in believing that p and believing that not p, and one can avoid this irrationality by dropping one of these beliefs, even if one thereby violates the verdicts of epistemic reason. We have something like the following:

Second Observation about Simple Consistency: One satisfies some norm of rationality by exiting the state of believing that p and believing that not p, even when one does not thereby satisfy the verdicts of epistemic reason on the question whether p.

Again Simple Consistency As Such seems to account for this in a straightforward way.

I believe, however, that there is another way to account for it. Whether or not we accept Simple Consistency As Such, the following (narrow-scope) rational requirement seems independently plausible:

Believed Epistemic Reason: If you believe at t that epistemic reason requires you to believe that p, then you are rationally required to form or sustain, going forward from t, the belief that p, and if you believe at t that epistemic reason requires you not to believe that p, then you are rationally required to revise or refrain from forming, going forward from t, the belief that p.⁵

Suppose that, in the case that we have been considering, you *believe* (falsely) that epistemic reason requires you not to believe that p and requires you to believe that not p. Then Believed Epistemic Reason rationally requires you not to believe that p and to believe that not p. This

⁵ Note that believing that epistemic reason requires you to believe that p, for the purposes of this requirement, need not involve saying anything to oneself *in foro interno*, let alone the words: “epistemic reason requires me to believe this.” If one says anything to oneself, one will say something like: “The evidence favors p” or “It’s most likely the case that p.”

explains why, in dropping the belief that p and continuing to believe not p, you satisfy a norm of rationality.

Now this would not explain the Second Observation as it stands. For all the Second Observation says, one satisfies some norm of rationality by exiting, in *any* way, a state of simple inconsistency. If, in the above case, one exited the state of simple inconsistency by continuing to believe that p and dropping the belief that not p, then one would *not* satisfy Believed Epistemic Reason, because that is not what one believes that epistemic reason requires.

But, on reflection, it is not clear that the Second Observation should be stated so permissively. It is not clear that one *would* satisfy some norm of rationality by defying one's own judgment about what one should believe. As it seems to one, the balance of evidence forbids believing that p and requires believing that not p, and yet one plunges ahead, in defiance of that apparent fact, and believes that p and does not believe that not p. That is irrational.⁶ If this is correct, then what we really have is the:

Second Observation about Simple Consistency, Revised: One satisfies some norm of rationality by exiting the state of believing that p and believing that not p, *in the way that one judges satisfies the verdicts of epistemic reason on the question whether p*, even when one does not thereby in fact satisfy the verdicts of epistemic reason on the question whether p.

And Believed Epistemic Reason explains this.

Something remains to be accounted for, however. Suppose, again, that epistemic reason requires one to believe that p and not to believe that not p and that one believes p and believes

⁶ One might say: "Yes, it is irrational, because it violates Believed Epistemic Reason. But, all the same, it satisfies a further norm of rationality, namely Consistency As Such." The question, however, is what evidence our ordinary practice supplies for positing this further norm of rationality.

not p, but, now, one simply doesn't believe that epistemic reason requires one to have a particular pattern of belief different from the pattern one has. One seems to be violating some norm in continuing to believe that p and to believe that not p, even when one does not believe that epistemic reason requires any change from one. Moreover, even if one came to a *false* judgment about the verdicts of epistemic reason, and then revised one's beliefs in accordance with this judgment, one would satisfy at least one of the norms that one was violating before forming that judgment.

Third Observation about Simple Consistency: One violates some norm in continuing to believe that p and to believe that not p, even when one does not believe that epistemic reason requires any change in those beliefs. Moreover, one can satisfy this norm without satisfying the verdicts of epistemic reason on the question whether p.

This norm cannot be Believed Epistemic Reason, since it applies even when one does not believe that epistemic reason requires any change in one's beliefs. And this norm cannot be a verdict of epistemic reason on the question whether p, since one can satisfy it by violating those verdicts. Again, the norm seems to be Simple Consistency As Such.

On reflection, however, it is not clear that Simple Consistency As Such explains all that is amiss in such a case. For all that Simple Consistency As Such says, one would satisfy it simply by losing one of the inconsistent beliefs. This fails to discriminate between the following two cases. In the first case, an electric shock simply eliminates the belief that p. One never realizes, of the contents one believes, that there is or was a conflict between them, nor would one have lost the belief that p but for the shock. In the second case, one realizes, of the contents one believes, that there is a conflict between them and then revises the belief that p accordingly. One wants to say that the subject gets something right in the second case that he does not get right in

the first case. He ought to realize, one wants to say, that there is a problem to be resolved. So, at very least, we would need to add to Simple Consistency As Such something to explain this “ought.”

But once we grant that at least part of what is amiss is that the subject fails to realize that there is a problem to be resolved, then it is no longer clear that we need to appeal to Simple Consistency As Such at all. If the subject believes that *p* and believes that not-*p*, then not only is he violating epistemic reason, but also he has compelling evidence that he is violating epistemic reason. Thus, he has, in most cases, conclusive reason to believe that epistemic reason requires him to change his beliefs and to deliberate about which change this is.⁷ This is the norm violated in the Third Observation. If, in response to this reason, he arrives at a belief (possibly mistaken) about which change epistemic reason requires, then he returns to the situation that we were considering earlier. If, at this point, he continues to believe *p* and to believe not *p*, then he is violating Believed Epistemic Reason.

What if, in response to this reason, the believer does not arrive at a belief about which change epistemic reason requires? He realizes that epistemic reason requires some change, but he remains undecided about which change this is. I am inclined to think that he is rationally required to suspend judgment. A doubt has been raised that epistemic reason permits him to believe that *p*, and he is presently deliberating whether it does. To believe that *p*, while still in the midst of this deliberation, without yet having concluded that he is permitted to believe that *p*, is to defy his own assessment of his reasons. My thought is that we should add to Believed Epistemic Reason a further clause:

⁷ Fabrizio’s point that this is also a requirement of rationality.

Believed Epistemic Reason, Revised: If you believe at t that epistemic reason requires you to believe that p , then you are rationally required to form or sustain, going forward from t , the belief that p , and if you believe at t that epistemic reason does not permit you to believe that p , *or if you are deliberating at t , in response to an active doubt, whether epistemic reason permits you to believe that p , but have not yet concluded that it does,* then you are rationally required to revise or refrain from forming, going forward from t , the belief that p .

[Mike T. raised some good objections to this, which I need to think through further.]

2. *Simple consistency: the negative thread*

So far, I have been working on the positive thread: trying to explain the appeal of Simple Consistency As Such in terms of other norms, which are not themselves norms of pure coherence. Now I turn to the negative thread: arguing that following Simple Consistency As Such is not itself a way of following the evidence toward truth and away from falsity.

The basic problem is this. We follow the evidence toward truth and away from falsity simply by believing as epistemic reason requires. Simple Consistency As Such does not require us to believe what epistemic reason requires. It is satisfied equally well if we have any other coherent pattern of belief. When we form a coherent pattern that is *not* what epistemic reason requires, in what sense are we following the evidence?

Simple Consistency As Such might be understood as a back-up, or second-best principle. By exchanging the incoherent pattern for some coherent pattern, one may not believe what epistemic reason requires. But one at least comes closer to believing what epistemic reason requires. The difficulty is that this is true only when the evidential situation is such that epistemic reason requires one neither to believe that p nor to believe that not p . In such a case,

any consistent alternative will be an epistemic improvement. At very least, one will lose one of the beliefs that epistemic reason requires one to lose. But when the evidential situation is such that epistemic reason requires one to retain one of these beliefs and to lose the other, then it is not true that each alternative is an epistemic improvement. One of the alternatives is to lose the belief that one is required to retain and to retain the belief that one is required to lose, and this is epistemically worse than staying put: retaining the belief that one is required to lose and retaining the belief that one is required to retain.

Next, it might be suggested that believing what epistemic reason requires is a complex achievement, achieved by taking several distinct steps. One of those steps is having a coherent pattern of belief. When we satisfy Simple Consistency As Such, therefore, we do at least part of what we must do to follow the evidence. We were never meant to follow SCAS alone. We were always meant to follow it along with other principles.

First, even on this suggestion, we do not follow the evidence by satisfying Simple Consistency As Such unless we also take the other necessary steps. Yet Simple Consistency As Such is supposed to be normative even when we do not take the other necessary steps. When we satisfy SCAS without taking the other steps, in what sense are we following the evidence?

Second, in what sense is satisfying SCAS even taking a “step” toward following the evidence? Suppose that epistemic reason requires one to believe that *p* and not to believe that not-*p*, but one conforms to SCAS by believing that not-*p* and not believing that *p*. In what sense has one done “part” of what epistemic reason requires? One might say: “Satisfying the disjunction—either believe that *p* and do not believe that not-*p*, or do not believe that *p* and believe that not-*p*, or do not believe that *p* and do not believe that not-*p*—is a *necessary condition* of believing what epistemic reason requires, namely: believing that *p* and not believing

that not-p. So when one satisfies the disjunction by believing that not-p and not believing that p, one takes a step toward believing what epistemic reason requires.” The worry is that, by the same logic, when one satisfies the disjunction—either believe that p and do not believe that not-p, or dance the Hokey Pokey—by dancing the Hokey Pokey, one takes a step toward believing what epistemic reason requires. [Other possibilities.]

It might seem promising to concede that complying with Simple Consistency As Such may leave one epistemically worse off in any particular case, but to contend that it tends to leave one epistemically better off in the long run. As I have tried to argue elsewhere, this is not true, at least not in way that might help Simple Consistency As Such.⁸ And even if it were true, Simple Consistency As Such is supposed to apply in each particular case. When we say that someone ought not both believe that p and believe that not p, we are saying something about what he ought to do here and now. This problem is parallel to the one that undoes traditional forms of rule utilitarianism. Why is there reason to follow, in this particular case, the rule that promises utility over the long run, if violating it, in this particular case, promises even more utility? Likewise, why is there reason to comply, in this particular case, with Simple Consistency As Such, which—we are supposing, contrary to fact—promises justified beliefs over the long run, if violating it, in this particular case, promises an additional justified belief?

[There seems no place for norms of pure coherence within the deliberative standpoint. Might norms of pure coherence have a role to play outside of the deliberative standpoint, as evaluative standards: say, standards of proper functioning? Not clear why they would even count as that.]

3. How is Believed Epistemic Reason normative?

⁸ “Why Be Disposed To Be Coherent?”

It might be argued that these arguments against SCAS prove too much. For they also tell against Believed Epistemic Reason, which I invoke to explain the Second Observation. If one falsely believes that epistemic reason requires one to believe that p, then Believed Epistemic Reason will require something that epistemic reason does not. If the only norms of belief are the norms of epistemic reason, then BER cannot be normative.

This is correct. Believed Epistemic Reason is not normative. Nevertheless, we can explain why Believed Epistemic Reason seems normative: why it seems, “from the inside,” when Believed Epistemic Reason applies to one, that one ought to believe what it requires, and why we can give something like advice, “from the outside,” to someone to whom (say) Believed Epistemic Reason applies, that takes the form of saying, “You ought to believe that p; it would be irrational of you not to.” In other words, we can make sense of our ordinary practice of treating Believed Epistemic Reason as though it were normative.

Let us approach this “Transparency Account,” as I have called it,⁹ by first considering matters “from the outside.” When we seem to advise someone who violates Believed Epistemic Reason that he ought rationally to believe that p, or that it would be irrational of him not to believe that p, I suggest, we are simply pointing out that he satisfies the antecedent of Believed Epistemic Reason. We are telling him, as we might put it, that *from his point of view, or as it seems to him*, the evidence says that he ought to believe that p. Thus, when we tell him that he ought rationally to believe that p, we are not ourselves conceding that epistemic reason requires him to believe that p. How, then, are we advising him to believe that p? By drawing his attention to the content of *his own judgment* that he ought to believe that p. Thus, a second-

⁹ “Why Be Rational?” *Mind* (2005).

person charge of irrationality, “But you ought to believe it! It would be irrational of you not to!” says, in effect: “Look: as it seems to you, you ought to believe it!”

Why does the advisee experience this as advice? More generally, why does being subject to Believed Epistemic Reason feel normative “from the inside”? Because, as things seem to the subject, she ought to believe that *p*. Such is the content of her own judgment, to which we are simply drawing her attention. In other words, given what the antecedent of Believed Epistemic Reason is, it will always seem to someone to whom BER applies that she ought to comply with it. We can explain why BER seems normative, while keeping faith with the idea that from the first-person standpoint of deliberation, our concern is simply to follow the evidence toward the true and away from the false.

This Transparency Account, however, cannot be given for SCAS. SCAS, unlike BER does not govern the relation between a subject’s judgment that he ought to have (or lack) an attitude and his having (or lacking) that attitude. So we are still left without an account of SCAS’s normative force, real or apparent: an explanation of why we should even be inclined to say that one “ought” to comply with SCAS.

4. Objections to the conception of epistemic reason

In trying to explain the appeal of Simple Consistency As Such, I have been assuming:

Stronger Evidence: Epistemic reason permits one to believe something only when the evidence that it is the case is stronger than the evidence that it is not the case.

One might wonder, however, what Stronger Evidence itself has to do with believing what is likely to be true and avoiding believing what is likely to be false. The answer, if there is one, is that it reflects the fact that avoiding falsity takes priority over acquiring truth. Epistemic reason

is inherently conservative, so that it permits the belief that p only when the chance of thereby acquiring a truth is greater than that of avoiding a falsehood.

It may help at this point to represent epistemic reason as though it were a kind of decision theory. This is an idealization, or stylization, of course, and as such, it risks some distortion. At the same time, however, it allows us to represent, in a particularly clear way, an exclusive concern with acquiring truth and avoiding falsity. Think of “decisions” as patterns of belief (which, of course, differ from genuine decisions in not being under one’s voluntary control). The “outcomes” are epistemically possible patterns of true and false propositions. The “payoff” of a given decision in an given outcome is a function that increases in beliefs in that decision that are true in that outcome and decreases in beliefs in that decision that are false in that outcome. The increase in the payoff that results from believing truly that p as opposed to not believing that p be some $T > 0$, and the increase in the payoff that results from not believing that p as opposed to believing falsely that p be some $F > 0$. The “probability” that p is the evidential support that p , $0 \leq E(p) \leq 1$. On the one hand, unless we suppose that degrees of evidence conform to the probability axioms, there will be cases in which there is no unique answer to the question whether we ought to believe that p . On the other hand, it may seem implausible that degrees of evidence conform to the probability axioms. In particular, it may seem implausible that for any two propositions, the evidence for each stand in some precise ratio. Isn’t there such a thing as *not* having evidence, or having insufficient evidence in *absolute* terms, for a proposition, which is something different from having insufficient evidence, in *comparative* terms, for the proposition: that is, having clear indications that the proposition is no more likely to be true than false? By way of a partial resolution of this dilemma, we might adopt the pretense that, for every proposition, there is a definite degree of evidence that behaves probabilistically, but for some

propositions, its degree of evidence is “unknown for the purposes of the epistemic decision.” Suppose, for example, that q_1, q_2, \dots, q_n is a partition of p , but that the only “known” degrees of evidence are $E(q_1)=.1$ and $E(q_2)=.1$. Then all that is “known” about $E(p)$ is that $E(p) \geq .2$. The epistemic decision whether to believe that p must be made on this basis.

What decision rule guides epistemic reason? The natural suggestion is maximizing expected value. Suppose for the moment that all relevant degrees of evidence are “known.” Then insofar as epistemic reason maximizes expected value, it permits believing that p only if:

$$E(p \text{ is true}) * T - E(p \text{ is false}) * F \geq 0,$$

or (barring cases in which $E(p \text{ is false})=0$) only if:

$$E(p \text{ is true})/E(p \text{ is false}) \geq F/T.$$

We get Stronger Evidence with two further steps. The first is to suppose that $E(p \text{ is false})=E(\text{not-}p)$. The second is to suppose that $F > T$: that, from the standpoint of epistemic reason, it is worse to believe the false answer to the question whether p than to fail to believe the true answer. If so, then epistemic reason permits the belief that p only if:

$$E(p) > E(\text{not } p),$$

which is just Stronger Evidence.

Now, $F > T$ has at least some independent support. It offers an explanation of:

Comparative Suspension of Belief: When the evidence that p and that not p are evenly balanced, one is required neither to believe that p nor to believe that not p .

One might think that, so long as one accepts Simple Consistency As Such, one could explain Comparative Suspension while denying $F > T$. But this is not so. If $F \leq T$, then when $E(p)=E(\text{not } p)$, epistemic reason permits believing p and permits believing not p . While Simple Consistency

As Such rules out adopting *both* beliefs, it does not rule out adopting only *one* of the two beliefs arbitrarily. $F > T$ also offers an explanation of:

Absolute of Suspension of Belief. When the evidence that p and that not p is insufficient in an absolute sense to decide the question, one is required neither to believe that p nor to believe that not p .

Here we relax our assumption that all evidence is “known.” Within the decision-theoretic model, there is absolutely sufficient evidence iff it is “known” whether $E(p)/E(\text{not-}p) \geq F/T$ (which may be known even if precise values for $E(p)$ and $E(\text{not-}p)$ are not). When the evidence is not absolutely sufficient, however, then we cannot decide whether to believe that p on the basis of maximizing expected value. We face, as it were, a decision under “uncertainty,” or unknown probabilities, rather than one under “risk,” or known probabilities. One natural suggestion is that we decide by assuming that $E(p) = E(\text{not-}p)$ and maximizing expected value on that assumption. Another suggestion is that we decide on the basis of maximin. Either of these natural decision rules leads us to suspend belief if and only if $F > T$.¹⁰

Nevertheless, one might protest that at least in certain contexts or, in some sense, “counterfactually,” $F \leq T$. [E.g. Popper.] In such cases, epistemic reason does not explain the First Observation. We must assume Simple Consistency As Such in order to explain it. To be

¹⁰ This view of degrees of evidence may thus avoid an objection that John Pollock and Joseph Cruz, *Contemporary Theories of Knowledge*, level against probabilism. If one has no information about p or q , other than that they are not logically related, then, intuitively, one should suspend judgment on p and on p or q . The probabilistic equivalent to suspending judgment is assigning a .5 degree of belief to each. Yet the probability axioms forbid this, Pollock and Cruz observe, since they require assigning a higher degree of belief to $(p$ or $q)$ than to p . It might seem that, likewise, if we assume that evidence is probabilistic, then we cannot justify suspending judgment on p and $(p$ or $q)$. But this is not so. While it follows that $E(p$ or $q) > E(p)$, it does not follow that $E(p) = .5$. Indeed, given that we have no further information about these propositions, it presumably is not “known” whether $E(p$ or $q)/E(\text{not } (p$ or $q)) \geq F/T$, or whether $E(p)/E(\text{not-}p) \geq F/T$. So we ought to suspend judgment on both p and $(p$ or $q)$.

clear about the dialectical context: While I find $F > T$ plausible, in part because it explains Comparative and Absolute Suspension, perhaps $F > T$ and Comparative and Absolute Suspension might fail in some contexts or counterfactually. In that case, I suppose, the First Observation would also fail. I do not claim that the First Observation need be true in all contexts or “counterfactually.” The positive thread requires that it be true in *enough* contexts, or as things *actually* are, to provide a theory of error for Simple Consistency As Such. But nothing more. Still, someone might insist that it sometimes is, or at least could be, the case that $F \leq T$, and even so, the First Observation does, or would, hold. Since epistemic reason cannot explain why, we need to assume SCAS.

Of course, this does nothing to explain how, if deliberation about what to believe is exclusively evidential, SCAS can be normative. Indeed, the assumption that $F \leq T$ only heightens the mystery. If $F = T$, if acquiring truth is just as important as avoiding falsity, then in at least some cases, believing both p and not- p will not be an unreasonable way to pursue truth and avoid falsity. And $F > T$, if acquiring truth is more important than avoiding falsity, then in at least some cases, believing both p and not- p will be the only reasonable way to pursue truth and avoid falsity. Nevertheless, one might argue, even if questions remain about the normativity of SCAS, these are questions that we must answer. We cannot reject SCAS.

While I do not have any decisive reply, let me point out two further difficulties with this suggestion. The first difficulty, the *Problem of Conflict*, arises, specifically, from the joint claim that SCAS is a valid principle of rationality and that $F < T$. If $F < T$, then epistemic reason will sometimes *require* us to *violate* SCAS. Whatever our view of the content of the verdicts of reason and the requirements of rationality, we might have expected that their contents were compatible, in the following senses:

Knowing Compatibility: It is not irrational to do what one knows reason requires.

Ideal Compatibility: It is never the case that, no matter what an agent might have done differently, he could not have avoided violating either the requirements of reason or the requirements of rationality. That is: If a subject is required at t by reason to X and required at t by rationality not to X , then there should be some earlier time t' and response R such that if the subject had given at t' response R , then he would not be both required at t by reason to X and required at t by rationality not to X .

Granted, there may be cases in which it is irrational of one to do what reason requires, because one *falsely believes* that reason prohibits it. Believed Epistemic Reason says as much. But if one *knows* what reason requires, then it should not be irrational of one to do it. Likewise, if one comes to the false belief that reason prohibits what it in fact requires, then, because of one's *past mistake*, it may be irrational of one to do what reason requires. But it shouldn't be fated in advance that there was simply no way to have avoided this conflict between reason and rationality. I do not mean to claim that all requirements must be compatible in these senses. Perhaps there could be a legal system in which one law requires you knowingly to violate another, or in which one is fated, no matter what one does, to violate some of its laws. What is more puzzling is that the two seemingly most fundamental, "natural" domains of requirement should be ineluctably at odds with one another.

"The Problem of Conflict does not arise," it might be replied, "if we assume that $F=T$. Then epistemic reason never requires us to believe that p and believe that not- p . Instead, when and only when $E(p)=E(\text{not-}p)$, it merely permits us to believe that p and merely permits us to believe that not- p . We need SCAS to explain why, in this special case, we are *not* permitted, in this case, to believe that p and believe that not- p ." It is unclear, however, whether epistemic

reason ever does merely permit us to believe that p. We do not seem to experience, in our doxastic deliberation, the kind of arbitrary liberty that the recognition of a mere epistemic permission would provide. This is the second difficulty: the *Problem of Permissiveness*.¹¹ To be faithful to the phenomenology of epistemic deliberation, we have reason to accept:

Epistemic Strictness: Epistemic reason never permits, but does not require, a belief.¹²

(In our decision-theoretic model: one is permitted to believe that p only when $E(p)/E(\text{not } p) > F/T$.)

Epistemic strictness may seem a stronger doctrine than it in fact is. It is compatible with mere permissions to believe that arise from what I go on to call “interest conditionality.” If there is interest conditionality, then there are two questions that bear on what, if anything, to believe regarding the question whether p. The first is *whether* to make up one’s mind whether p. This is not a question decided by epistemic reason. It is answered by balancing the benefits of making up one’s mind, such as the instrumental or intrinsic importance of the knowledge that might thereby be gained, against the costs, such as the time and energy that might be expended on making up one’s mind. The second question is *how* to make up one’s mind, if one makes it up: what to believe, if one is to have an opinion on the question. This is decided by epistemic reason, guided by the aims of avoiding falsehood and acquiring truth. If we accept epistemic strictness along with interest conditionality, then one is always required to make up one’s mind in a particular way, *if* one makes it up at all. But this does not mean that one is required to make

¹¹ It is no use to reply: “The reason why we never experience this kind of freedom is that, when epistemic reason permits believing p and believing not-p, SCAS forbids it.” Even with SCAS, we would still have this freedom. We would be permitted to believe p and not believe not-p, and permitted to believe not-p and not believe p, and permitted neither to believe p, nor to believe not p.

¹² Compare Roger White, “Epistemic Permissiveness,” *Philosophical Perspectives* (forthcoming).

up one's mind in that way. One may be permitted not to make up one's mind at all. In such a case, if one does not make up one's mind, then one is not violating epistemic reason in failing to believe what epistemic reason requires one to believe, if one considers the question. This kind of mere permission does not reintroduce the kind of freedom to believe arbitrarily that epistemic strictness is meant to rule out. The liberty in question is simply the liberty not to make up one's mind on a question. If one forgoes that liberty, and turns to making up one's mind, then one has no liberty in how one does so. That is determined by the evidence, at least if Epistemic Strictness holds. And if Epistemic Strictness holds, then the gap never opens up that SCAS is invoked to close.

Now consider a different objection: "The evidence for one that p is a function of the beliefs that one has. Thus to say, for example, that the balance of evidence that p over evidence that not p requires one to believe that p is to invoke a norm of coherence. One is required to believe p because it coheres with one's other beliefs. Thus, the attempt to explain norms of coherence in terms of norms of evidence presupposes what it attempts to explain."

If there is some "norm of coherence" presupposed by evidence, this is not a norm of *pure* coherence. No strongly wide-scope norm can make sense of the idea that believing that p is supported over not believing it, since any wide-scope norm is satisfied equally well by either alternative. What is possible is that the concept of evidence presupposes a norm of coherence of another kind. It might be suggested, for example, that when we evaluate the evidence that p, we hold your other beliefs fixed. There is evidence for you that p just to the extent that your other beliefs "support," in some sense, p.

I don't think that this conception of evidence would change much. There would still be good reason to accept Stronger Evidence, now read as: "One has sufficient reason to believe that

p only if one's other beliefs support p more strongly than they support not-p." And if Stronger Evidence were accepted, then we could offer the same explanation of the First Observation. It might seem that if we identify evidence with what one believes, then we must assume a norm of pure coherence. Otherwise, the evidence may be inconsistent. But, on any view, the evidence may be "inconsistent," in the sense that there can be some evidence that p and some evidence that not-p.

It is a further question whether it makes sense, within the evidentialist view that animates this paper, to accept the suggested conception of evidence:

Beliefs Provide Evidence: There is evidence for you that p just to the extent that your other beliefs support p.

We can begin by asking whether BPE holds that the evidence *is* that you have those other beliefs. Such a view would hard to be reconcile with the idea that evidence that p "indicates," "suggests," or "makes it more likely," that p. For it seems more plausible that it is what you believe, and not the fact that you believe it, that makes it likely that p. Somewhat more precisely:

Nonpsychologism about evidence: There is no basic evidential principle to the effect that my having some belief, or some other attitude for which I can have or lack reasons, is evidence for me that something is the case.

Nonpsychologism does not deny basic principles to the effect that my having a mental state for which I cannot have or lack reasons, such as a perceptual experience, can be evidence. Nor does Nonpsychologism deny that my having a belief (or other attitude for which I can have or lack reasons) can be evidence. After all, that I believe that p is evidence that someone believes that p, that I hold some paranoid belief is evidence that my medication is wearing off, and that I intuitively believe—that I have a "hunch"—that this is the right play at this point in this game is

evidence, given that my hunches are usually correct, that this is the right play. In such cases, the evidential significance of my belief can be explained by a *more basic* evidentiary principle that makes no reference to my attitudes.¹³ That I believe that p is evidence that someone believes that p in the way in which Fa is evidence that there exists an x such that Fx. That I hold some paranoid belief is evidence that my medication is wearing off in the way in which the obtaining of an effect is evidence of the obtaining of a likely cause. That I have a hunch that this is the right play is evidence that it is the right play in the way in which the judgment of a reliable authority that p is evidence that p. One indication that the underlying evidentiary principle does not refer to *my* attitudes is that someone else's belief can be evidence for me in precisely the same way.

However, a proponent of BPE need not deny Nonpsychologism. He might agree that it is the *contents* of your other beliefs that are the evidence, or make it more likely, that p. Indeed, it might be said, that is just what it means to say that those beliefs “support” p. If we accept Nonpsychologism, then it seems hard to deny:

Factivity: That q is evidence for you that p only if it is the case that q.

Suppose that I am a reflective, knowledgeable believer, who deliberates about what to believe on the basis of the evidence, but also has the correct view, whatever it may be, about what evidence is. When I, loosely speaking, take q to be evidence that p, or to make it more likely that p, what, more exactly speaking, do I take to be evidence, or to make it more likely, that p? Not the fact that I believe that q, if I accept Nonpsychologism. Nor the fact that q, if I deny Factivity. So what exactly? Apparently, the content of my belief, whether or not it is true, that q. But it is hard to know what it might mean to say that the content of my belief that q, whether or not it is

¹³ One might be tempted to say that in such cases, my belief only contingently has this evidential significance. But my belief that p is not only contingently evidence that someone believes that p.

true, makes it more likely that p. It isn't the bare proposition that q that makes it more likely that p. Otherwise we run the risk of there being maximal evidence for and against everything. Nor is it my believing this proposition, if, again, I accept Nonpsychologism. And whatever it means, accepting it would seem to commit me to the view that if it were *false* that q (at least so long as I continued to believe it), it would still be *just as likely* that p. But when I take q, loosely speaking, to be evidence that p, I do not believe that if it were false that q, it would still be as likely that p (again, so long as I accept Nonpsychologism).

If BPE is to be compatible with Factivity, then it must be understood as:

True Beliefs Provide Evidence: There is evidence for you that p just to the extent that your other true beliefs support p.

And TBPE is even further from a norm of pure coherence. Not only is it (or the requirement of epistemic reason that it supports) narrow scope, but also it is not a norm of rationality. And once we accept Factivity, one might wonder, why even accept TBPE? Why not say that facts that are not believed can be evidence?

The answer may be as follows: Evidence must be limited to what is, in some sense, available to the believer. Otherwise, talk of what the evidence suggests is true loses its point. Why not simply talk of what is, in fact, true? And evidence is available to the believer only insofar as he believes it. But is it obvious that this is the only viable conception of availability? It might be suggested, for example, that q is available if you have some nondoxastic perceptual experience that q. More ambitiously, it might be suggested that q can be available even if do not actually have some mental state that q. For example, q might be available to you because you are in a position to become aware of it by performing some mental act: because, e.g., it is in your field of vision and you would see it if you only paid attention. Still more ambitiously, it might

be suggested that q can be available to you because you are in a position to become aware of it by performing some physical act: because, e.g., it would be in your field of vision if you would only look. Denying that that which requires a mental act can be available seems potentially overly restrictive. While we certainly come to believe much of what we believe without performing mental acts, this is not always so. The difficult, and interesting, questions often enough require that we focus our attention, pose questions to ourselves, imagine how hypotheses would play out, visualize, and so on. And if we accept that what requires a mental act can be available, then we may begin to wonder why what requires a physical act cannot. After all, we cannot do all that much in our heads. Sometimes, we need to put things down on paper, for example. Why can't the truths that we might, or might have, reached in these ways count as available to us?

5. *Simple closure*

Thus far we have been considering simple consistency. Now let us turn to simple closure. To a first approximation, at least, we have the:

First Observation about Simple Closure: If q is a logical consequence of p , and one believes that p , but does not believe that q , then one violates some norm.

Again, it may seem natural to suppose that this observation is accounted for by:

Simple Closure As Such: If q is a logical consequence of p , then one is rationally required (either not to believe p , or to believe q).

And, again, my suggestion is that we can explain this fact instead by considering what epistemic reason requires.

When q is a logical consequence of p , any evidentiary argument that p —any argument that suggests that p is the case—is an evidentiary argument that q . For if p is true, then q must be true. So we get:

Transmission: If q is a logical consequence of p , then the evidence that q is at least as strong as the evidence that p .

By Transmission, if q is a logical consequence of p , then $E(q) \geq E(p)$. By contraposition, $E(q)/E(\text{not } q) \geq E(p)/E(\text{not } p)$. If we continue to assume Epistemic Strictness, then epistemic reason permits p if and only if $E(p)/E(\text{not } p) > F/T$ and $E(p)$ is absolutely sufficient. And this is true if and only if $E(q)/E(\text{not } q) > F/T$ and $E(q)$ is absolutely sufficient. Thus, if epistemic reason permits p , then it requires q . If one believes that p , but fails to believe that q , then one is violating some verdict of epistemic reason.

Put another way, there are three possible evidential situations:

(1) epistemic reason forbids p and forbids q ,

(2) epistemic reason forbids p and requires q ,

and

(3) epistemic reason requires p and requires q .

In the first situation, one violates the first verdict if one believes that p but does not believe that q . In the second situation, one violates both verdicts. In the last situation, one violates the second verdict. In each of these evidential situations, therefore, if one believes that p but does not believe that q , then one violates some verdict of epistemic reason. This accounts for the First Observation.

There are Second and Third Observations about Simple Closure that parallel the Second and Third Observations about Simple Consistency. And they can be explained in parallel ways:

by appeal to Believed Epistemic Reason, and by appeal to the fact that when one has simply unclosed beliefs, one has good evidence that epistemic reason requires some change from one.

As before, one might urge that, while the First Observation always holds, there are cases in which Transmission cannot explain it. We must assume Simple Closure As Such in order to explain the First Observation in these cases. One potential opening for Simple Closure As Such arises from what we might call “epistemic conditionality.” One might suggest, plausibly, that Transmission should be restricted to logical consequences that the relevant believer is in a position to know. Do I necessarily have as strong evidence for the Axiom of Choice, say, as I have for Zorn’s Lemma, if I am not in a position to see that the former follows from the latter? If Transmission is restricted in this way, then there will be cases in which epistemic reason requires one to believe that p, but does not require one to believe that q, because one is not in a position to know that q is a logical consequence of p. One might then suggest that we need Simple Closure As Such to explain why the First Observation holds in these cases.

However, this seems an unstable position. If one believes that Transmission is epistemically conditional, then one seems pressed to accept that the First Observation and Simple Closure As Such are also epistemically conditional. If it seems unfair to say that epistemic reason requires you to draw the consequence, because can’t be expected to know that it is a consequence, then it seems just as unfair to say that it would be irrational of you to fail to draw the consequence, because you can’t be expected to know that it is a consequence. The two views seem to stand or fall together.¹⁴

¹⁴ It is not clear obvious how a defender of Simple Closure As Such can introduce an epistemic condition. One possibility is to introduce it as an “objective” condition, so that we have:

Simple Closure with Objective Epistemic Condition: If q is a logical consequence of p that one is in a position to know, then one is rationally required (either not to believe p, or to believe q).

A second kind of conditionality, “interest conditionality,” has to do with the importance to the subject of having an opinion about what follows from the contents of his beliefs. It might be said that if the subject does not care, or alternatively, if the subject has no reason to care, about having an opinion whether p, then he is not prohibited from believing contents from which q follows, while not believing q. If there is interest conditionality, then there may be cases in which:

- (i) epistemic reason requires one to believe that p if one makes up one’s mind whether p,
- (ii) epistemic reason requires one to believe that q if one makes up one’s mind whether q,
- (iii) one is required to, or does, make up one’s mind whether p,
- (iv) and so is required to believe that p,
- (v) but one is not required to make up one’s mind whether q, and does not make up one’s mind whether q,
- (vi) and so is not required to believe that q.

It might then be argued that we need Simple Closure to explain why the First Observation holds in such a case.

My reply is very similar to my reply to the earlier point about epistemic conditionality. The considerations that might lead one to accept that epistemic reason is interest conditional would seem equally to lead one to accept that the First Observation is itself interest conditional.

But this introduces a fact about epistemic reason into the rational requirement, so that it is no longer a “pure” rational requirement. [Broome on “obvious” consequence.] The other possibility is to introduce it as one of the “subjective” disjuncts, so that we have:

Simple Closure with Subjective Epistemic Condition: If q is a logical consequence of p, then one is rationally required (either not to believe p, or not to believe that q is a logical consequence of p, or to believe q).

But then, it seems worth noting, Simple Closure As Such is vulnerable to complaints that are often leveled against Believed Epistemic Reason, that it is, by itself, powerless to criticize those who hold false logical beliefs.

First Observation about Simple Closure with Interest Condition: If q is a logical consequence of p , and one is required to, or does, make up one's mind whether q , and one believes that p , but does not believe that q , then one violates some norm.

Simple Closure would not be needed to explain this fact.¹⁵

One might argue, next, that different questions can have different evidential requirements. What counts as sufficient evidence for one to believe that p or to believe that not p may be weaker than what counts as sufficient evidence for one to believe that q or to believe that not q . Although we have been assuming that F/T is constant for all questions, it might be the case that it depends questions, and, in particular, F/T for the question whether $p < F/T$ for the question whether q . It might be more important not to have a false belief whether q than not to have a false belief whether p . If so, then the most that follows from the fact that epistemic reason permits one to believe that p is that $E(q)/E(\text{not}q) \geq E(p)/E(\text{not}p) > F/T$ for the question whether p . It might be the case that $E(q)/E(\text{not}q) < F/T$ for the question whether q , in which case epistemic reason would not permit one to believe that q .¹⁶ Notice that in this case epistemic reason *requires* one to believe that p but does *not permit* one to believe that q . Assuming Simple Closure As Such would then lead to the Problem of Conflict.

¹⁵ If one accepts this revision, then one would presumably also revise Simple Closure to include a "subjective" interest condition, such as:

Simple Closure with Interest Condition: If q is a logical consequence of p , then one is rationally required (either not to believe p , or not to believe that one is required to make up one's mind whether q , or to believe q).

It is true that epistemic reason cannot, whereas Simple Closure with Interest Condition can, explain what is irrational about believing that p , *falsely believing* that one is required to consider whether q , and failing to believe that q . However, Believed Epistemic Reason can explain this. It is irrational either because the person believes that epistemic reason requires him not to believe that p , or because the person believes that epistemic reason requires him to believe that q (since he believes that he is required to make up his mind on the question).

¹⁶ The First Observation might fail in a similar way if epistemic reason were "risk-averse": if F increased with each additional false belief faster than T increased with each missed true belief.

Finally, one might deny Epistemic Strictness, despite the Problem of Permissiveness. Then, in the special case in which $E(q)/E(\text{not } q)=E(p)/E(\text{not } p)=F/T$, epistemic reason would merely permit believing p and merely permit believing q . In this case, one would not violate epistemic reason by believing that p , but not believing that q . So one might argue that we need Simple Closure As Such to explain the First Observation.

6. *Are consistency and closure, in general, required of us?*

Having considered the cases of simple consistency and closure, let us now turn to their generalizations. Some will accept the following:

First Observation about General Consistency: If not p_{n+1} is a logical consequence of p_1, p_2, \dots, p_n , then if one believes p_1 and believes p_2 and ... and believes p_{n+1} , then one violates some norm.

First Observation about General Closure: If p_{n+1} is a logical consequence of p_1, p_2, \dots, p_n , then if one believes p_1 and believes p_2 and ... and believes p_n , but does not believe p_{n+1} , then one violates some norm.

Call these, together, the First Observations about General Logical Coherence. Again, it may seem plausible that the norms in question are as-such rational requirements:

General Consistency As Such: If not p_{n+1} is a logical consequence of p_1, p_2, \dots, p_n , then one is rationally required (either not to believe p_1 , or not to believe p_2, \dots , or not to believe p_{n+1}).

General Closure As Such: If p_{n+1} is a logical consequence of p_1, p_2, \dots, p_n , then one is rationally required (either not to believe p_1 , or not to believe p_2, \dots , or not to believe p_n , or to believe p_{n+1}).

Can we instead explain the First Observations in terms of the verdicts of epistemic reason? I do not think so. The difficulty arises from patterns of evidence such that, while p_{n+1} is a logical consequence of p_1, p_2, \dots, p_n , it is nevertheless the case that $E(p_i)/E(\text{not } p_i) > F/T$, $1 \leq i \leq n$, and $E(\text{not } p_{n+1})/E(p_{n+1}) > F/T$. In such a case, epistemic reason requires one to believe that p_1 , epistemic reason requires one to believe that p_2, \dots , epistemic reason requires one to believe that p_n , epistemic reason requires one not to believe that p_{n+1} (violating General Closure As Such), and epistemic reason requires one to believe that not p_{n+1} (violating General Consistency As Such). I see no way of avoiding this result if we understand epistemic reason as a faculty that aims at truth and away from falsity in light of the evidence, and grant that evidence can assume such a pattern.¹⁷

Can the evidence assume such a pattern? The familiar paradoxes of the lottery and the preface present us with cases in which, intuitively, the evidence does just that. In the case of the lottery, there is, for each ticket, much stronger evidence that it will lose than that it will not lose, but also much stronger evidence that it is not the case that all of those tickets will lose than that it is the case that all of those tickets will lose. Thus, it seems reasonable to believe of each ticket that it will lose, while believing that one of the tickets will win. In the case of the preface, there is, for each claim made in a book, much stronger evidence that the claim is the case than that it is not the case, but also much stronger evidence that it is not the case that all of those claims are true than that it is the case that all of those claims are true. Thus, it seems reasonable for the author to hold, and to express in his preface, the belief that some of those claims are false.

The doubts that have been raised about these concrete cases strike me as unpersuasive. Some doubt that real-life authors actually believe at the same time each claim that they have

¹⁷ This is hardly an original point. For a much more thorough and resourceful discussion, see David Christensen, *Putting Logic in Its Place*.

written in their books, while also believing the negation of the long conjunction of those claims. Who can hold so much in his head at the same time? But why not simply stipulate that an author might have everything available, as we often stipulate unlikely contingencies (trolleys and suchlike) to test the implications of normative principles?¹⁸ And even if we could somehow dismiss these concrete illustrations, the deeper worry is that there is no principled way to rule out such patterns of evidence. While degrees of evidence may not always behave like probabilities, there seem to be at least special cases in which the relevant degrees of evidence do behave like probabilities: namely, when the relevant evidence consists in the available relative frequencies of events of the types in question.¹⁹

In sum, we cannot explain the Observations about General Logical Coherence by appeal to epistemic reason. Nor, I think, can we explain the Observations by appealing to General Consistency and Closure As Such. We cannot understand these norms of pure coherence as ways of pursuing the evidence.²⁰ Moreover, since epistemic reason sometimes *requires* the

¹⁸ A more interesting, but also less focused, doubt is that it seems somehow fishy that we accept inconsistency between a claim in the preface and claims in the body, but do not accept inconsistency among claims in the body. I suspect that this is because the body, unlike the combination of the preface and the body, is typically presented as a theory or a history: as a way the world *could* be or *could* have been. If there is an inconsistency among those claims, then it forfeits its status as that.

¹⁹ Some may suggest that if one's beliefs are incoherent in this way, then one is required to seek out more evidence on the basis of which one could hold coherent beliefs. This view seems to concede what is primarily at issue: that one is not under any requirement, here and now, to make one's beliefs coherent. In any event, it is not clear that one is always required to seek out more evidence. One might already have as much evidence as is ever reasonable to expect for a contingent proposition. And if one is required to seek out more evidence, this requirement seems independent of the incoherence in one's beliefs. Someone with the same beliefs except for the preface belief in his own fallibility, in the same situation, would seem to be under the same requirement to look for more evidence, even though his beliefs would not be incoherent.

²⁰ This follows immediately from the earlier finding that Simple Consistency and Closure As Such, which are special cases, are not ways of pursuing the evidence.

violation of General Consistency and Closure As Such, invoking them would lead to the Problem of Conflict.

7. Probabilistic coherence

Many probabilists agree that the Observations about General Logical Coherence cannot be accounted for. Indeed, some suggest that this fact shows that our subject ought to be not full beliefs, but instead degrees of belief. It makes no sense to claim that degrees of belief must conform to General Logical Coherence. Instead, they insist on the:

First Observation about Probabilistic Coherence: If one has degrees of belief that do not conform to the probability axioms, i.e.,

- (i) that for all degrees of belief that $p (=D(p))$, $0 \leq D(p) \leq 1$,
- (ii) where not q is a logical consequence of p , $D(p \text{ or } q) = D(p) + D(q)$, and
- (iii) where p is a logical truth, $d(p) = 1$;

then one violates some norm.

I ask two questions about this view. The first is whether epistemic reason can account for the First Observation about Probabilistic Coherence. The second is whether we should adopt a framework of degrees of belief specifically on the grounds that we cannot account for the Observations about General Logical Coherence within framework of full belief.

Perhaps because probabilistic coherence is not as intuitive as logical coherence, there have been more deliberate attempts to explain it. Some of these appeal to epistemic reasons for belief; others appeal to nonepistemic reasons. I focus on epistemic explanations. But, as it happens, some nonepistemic explanations are analogous in the respects that most interest us.

In “A Nonpragmatic Vindication of Probabilism,” James Joyce suggests that, whereas for full beliefs, the aims of epistemic reason are to avoid false beliefs and to acquire true ones, for

degrees of belief, the aim of epistemic reason is simply to avoid “inaccuracy.” On any plausible measure of “inaccuracy,” one’s set of degrees of belief is less inaccurate to the extent that one’s degrees of belief in truths are closer to one and one’s degrees of belief in falsehoods are closer to zero. Joyce then proves that any accuracy measure that meets certain constraints has the following property: For any set of degrees of belief that is probabilistically incoherent, there is some probabilistically coherent set that is less inaccurate, by that measure, no matter what the actual state of the world.

Joyce’s proof does make a compelling case for the following claim: that if the measure of accuracy that structures epistemic reason meets the constraints, then whatever set of degrees of belief epistemic reason requires will be a probabilistically coherent set. Yet this conclusion does not support a norm of pure coherence to have probabilistically coherent degrees of belief: i.e.,

Probabilism As Such: One is rationally required to have, for all propositions p , $D(p)$ such that:

- (i) $0 \leq D(p) \leq 1$,
- (ii) if not q is a logical consequence of p , $D(p \text{ or } q) = D(p) + D(q)$, and
- (iii) if p is a logical truth, $D(p) = 1$.

Some probabilistically coherent sets of degrees of belief are *not* guaranteed to be more accurate no matter what the state of the world. Satisfying Probabilism As Such by adopting one of these sets may leave one epistemically worse off. The situation is parallel to Simple Consistency As Such.²¹

²¹ The nonepistemic explanation of the First Observation, or at least one such explanation, assumes that one’s degree of belief in a proposition corresponds to one’s willingness to bet on its being true. It is then shown that probabilistically incoherent beliefs will lead one, in certain circumstances, to place bets that are guaranteed, no matter what the state of the world, to leave the aims that one seeks to achieve by those bets worse achieved. If the aims that one seeks to

Although Joyce's proof does not tell us how to determine which probabilistically coherent set epistemic reason requires, a natural suggestion is that epistemic reason requires the set that minimizes expected inaccuracy, given the evidence. Let $T(D(p))$ represent the contribution to inaccuracy of $D(p)$ when p is true, and $F(D(p))$ represent the contribution to inaccuracy of $D(p)$ when p is false. Then the expected increase in inaccuracy of degree of belief $D(p)$ is:

$$E(p)*T(D(p)) + E(\text{not-}p)*F(D(p)).$$

This is minimized when:

$$\textit{Minimum: } E(p)/E(\text{not } p) = - F'(D(p))/T'(D(p))^{22},$$

achieve are aims that one has reason to achieve, then probabilistically incoherent degrees of belief will lead one to achieve less of what one has reason to achieve. Probabilistically coherent beliefs, by contrast, will not lead one to place such bets. If we were willing to countenance nonepistemic reasons for belief, then we would have an explanation of the First Observation. This explanation does not support the pure norm, Probabilism As Such. The most that we may say is that if one has a probabilistically incoherent set, then there is *some* particular probabilistically coherent set that will lead one to better achieve one's aims, by plugging whatever openings one has left for a "Dutch book." Presumably, there is some *other* probabilistically coherent set (e.g., a set including a .99 degree of belief that one will not be harmed by ingesting some poison) that, while insulating one from Dutch books, can nonetheless be expected to be worse, in the given circumstances (e.g., in which the poison is all too easy to ingest), for the aims that one has reason to achieve than being exposed to Dutch books. Notice that there is nothing illicit about taking into account the facts of the matter (e.g., that, in fact, one will be gravely harmed by ingesting the toxin) in determining whether one should trade one's incoherent set of degrees of belief for this particular coherent set. One's pragmatic reasons depend on the facts of the matter. However (as Richard Bradley pointed out to me in conversation) the pragmatic, Dutch-book argument does give one at least a *pro tanto* reason to conform to Probabilism. Any set of probabilistically coherent degrees of belief at least avoids a Dutch Book, so there is at least that to be said for it, whatever else there is to be said against it. The same cannot be said for the nonpragmatic argument.

²² Although this is only a necessary condition of minimization. The treatment parallels Allan Gibbard's in "Valuing Truth," except that degrees of evidence replace what for Gibbard are the believer's initial degrees of belief.

where $F'(D(p))$ and $T'(D(p))$ are the first derivatives evaluated at $D(p)$. If we assume a “least squares” measure of inaccuracy, i.e., $T(D(p)) = [1-D(p)]^2$ and $F(D(p)) = D(p)^2$, which satisfies Joyce’s constraints, then we have a minimum where:

$$E(p)/E(\text{not } p) = D(p)/[1-D(p)]$$

If degrees of evidence conform to the probability axioms (so that, among other things $E(\text{not } p) = 1 - E(p)$), then inaccuracy is minimized when $E(p) = D(p)$. We should apportion degrees of belief in proportion to the evidence.²³

Our second question, recall, was whether, setting aside whatever other, perhaps entirely compelling, reasons we might have for embracing a framework of degrees of belief, we should do so specifically on grounds that we cannot account for the Observations about General Logical Coherence. Probabilism, it is often suggested, is the lesson of the preface and lottery paradoxes.²⁴ The reasoning seems to be this:

- (i) Whatever the proper conception of belief, it must be the case that belief, on that conception, ought to be generally coherent, according to some understanding of general coherence that is distinct from simply what the evidence supports.
- (ii) The only such conception of general coherence for full belief is general logical coherence.

²³ Now suppose, as before, that the degrees of evidence for some propositions are “unknown.” How should we distribute degrees of belief? We might suppose, first, that our degrees of belief should conform to whatever constraints degrees of evidence are known to satisfy. Suppose that it is known only that p entails not q , $E(p \text{ or } q) = .6$, and $E(p) \geq 2E(q)$. Then it should be the case that $D(p \text{ or } q) = D(p) + D(q) = .6$ and $D(p) \geq 2D(q)$. But this still leaves a range of possible values for $D(p)$ and $D(q)$. In the case of full belief, we suggested, first, assuming that p and not- p were equally supported. But to assume that $E(q) = .5$ would violate the known constraints. The other suggestion was to adopt maximin as our decision rule. In this case, this would mean selecting $D(p)$ and $D(q)$, subject to whatever constraints the evidence is known to satisfy, so as to minimize the maximum possible inaccuracy. We do this by selecting $D(p)$ and $D(q)$ so as to minimize $\max \{ |D(p) - .5|, |D(q) - .5| \}$: i.e., $D(p) = .4$ and $D(q) = .2$.

²⁴ Recently by David Christensen, *Putting Logic in Its Place*.

- (iii) At least one such conception of general coherence for degrees of belief is probabilistic coherence.
- (iv) But it is not true that full beliefs ought to be generally logically coherent.
- (v) It is true that degrees of belief ought to be probabilistically coherent.
- (vi) Therefore, the proper conception of belief is one of degrees of belief.

There are two problems with this line of thought. The first is that (i) is an oddly abstract assumption to make. It is one thing to find some *specific* conception of coherence compelling and assume that belief must be coherent according that conception. But what motivates the thought that belief must be coherent according some conception, we know not which?

The second problem is that (v) is not as secure as it first looks. What compels us to reject the Observations about General Logical Coherence are the aims that it seems natural to attribute to epistemic reason insofar as it traffics in full belief: the aims of avoiding falsehood and of acquiring truth. There are bound to be cases, such as the preface and the lottery, in which the pursuit of these aims requires incoherent beliefs. What Joyce has shown is that there are *some* plausible conceptions of the aim of epistemic reason insofar as it traffics in degrees of belief—accuracy measures that satisfy his constraints—that support the Observation about Probabilistic Coherence. The difficulty is that there are other, not obviously implausible conceptions of the aim of epistemic reason that do not support the First Observation about Probabilistic Coherence. Allan Gibbard suggests a measure that is more concerned with accuracy at the extremes, so that, for example, $T(.9) - T(.99) > T(.4) - T(.49)$. A simple candidate (no doubt too simple) would be $T(D(p)) = (1 - D(p))^{1/2}$ and $F(D(p)) = D(p)^{1/2}$. Substituting the relevant values in *Minimum*, even assuming that evidence is probabilistic, we get:

$$E(p)/[1 - E(p)] = [(1 - D(p))/D(p)]^{1/2}$$

Inaccuracy would not, in general, be minimized by probabilistically coherent degrees of belief.²⁵ Another possibility, which Gibbard does not mention, is a measure that is more concerned with minimizing the degree of belief in falsehoods than with maximizing the degree of belief in truths. This would be the partial-belief analog to the assumption that $F > T$. For example, if $T(D(p)) = (1 - D(p))^2/2$ and $F(D(p)) = D(p)^2$, then inaccuracy would be minimized where:

$$E(p)/[1 - E(p)] = 2D(p)/[1 - D(p)].$$

Again, inaccuracy would not, in general, be minimized by probabilistically coherent degrees of belief.²⁶ In sum, the problem is that the First Observation about Probabilistic Coherence is vulnerable to its own version of the preface paradox: a case in which aiming faithfully to represent reality, on a not implausible conception of what that amounts to, leads to degrees of belief that are not probabilistically coherent.

8. How is logic normative for belief?

It seems hard to deny that, in some sense, logic governs belief, that it supplies rules for thought. To the extent that the negative thread casts doubt on the norms of pure coherence, it casts doubt on a certain understanding of how logic governs belief: that logic governs belief directly, that our beliefs themselves are constrained to be logically consistent and closed. The positive thread, by contrast, reflects a different understanding of how logic governs belief. On this view, epistemic reason alone directly governs belief. Our beliefs are constrained simply to follow what the evidence suggests is true and to avoid what the evidence suggests is false. Logic governs belief indirectly, by structuring the evidence. Logic, so to speak, tells epistemic reason

²⁵ Let p entail not q , $E(p) = E(q) = 1/5$. Then inaccuracy is minimized when $D(p) = D(q) = 16/17$ and $D(p \text{ or } q) = 9/13$. This measure violates Joyce's constraints of Weak Convexity and Symmetry.

²⁶ Let $E(p) = 1/2$. Then inaccuracy is minimized when $D(p) = D(\text{not } p) = 1/3$. This measure violates (at least) Joyce's constraints of Dominance and Normality.

about possible patterns of truth and falsity. If q is a logical consequence of p , for example, then q is true, if p is true. So any evidential support that p has is also evidential support that q has. Thus, $E(q) \geq E(p)$. If $\text{not-}p$, then p is false, and conversely. So any evidential support that p is false is evidential support that $\text{not-}p$, and conversely. Thus, $E(p \text{ is false}) = E(\text{not-}p)$. Epistemic reason then takes these patterns into account in determining how best to pursue the aims of acquiring truth and avoiding falsity in light of the evidence.²⁷

This is not to say that logic constrains belief in virtue of determining what counts as conforming to the norms of probabilistic coherence, which in turn directly constrain belief.²⁸ First, what are being directly constrained are degrees of evidence, not degrees of belief. Those degrees of evidence, as Gibbard notes, may call for degrees of belief that are not probabilistically coherent. Second, we need not be concerned with degrees of belief at all. What may be formed in light of the evidence so constrained may be full beliefs. Finally, these relationships may hold even if evidence does not behave like probabilities in other respects.

Logic, therefore, is an “input” into epistemic reason. But this fact offers, on its face, no guarantee that the “outputs” of epistemic reason, the beliefs that it requires, will themselves conform to a logical pattern. Nor is there compelling reason to expect that it should conform to a probabilistic pattern, or to any other pattern, besides simply that of being sufficiently supported by the evidence. As it happens, and indeed somewhat surprisingly from this standpoint, the information that logic gives epistemic reason *does* guarantee (at least given some assumptions about Epistemic Strictness and the ratio of T to F) that the beliefs that epistemic reason requires

²⁷ Of course, there are other sources of information, besides logic, about possible patterns of truth and falsity. If logic governs belief in a special way, it may be because it is somehow specially available to us, and so its contribution to epistemic reason is specially insensitive to our particular evidential situation. Here we revisit the issues that we considered in connection with “epistemic conditionality.”

²⁸ This is where Christensen suggests logic should be “put.”

will be simply consistent and closed. But it does not guarantee, as far as I can see, that the beliefs that epistemic reason requires will be logically coherent in general.

The question with which we are then left is whether logic places a *further* constraint on our beliefs: whether it does double duty. Does logic not only inform our epistemic deliberations about what to believe on the basis of the evidence, but also constitute an additional constraint to be applied to what issues from those deliberations? Here, I think, we should answer “no.” It is clear why logic should matter to our beliefs in the first way: that is, by telling us what must be true or false when something else is true or false. But why should logic matter to our beliefs in the second way: as an independent constraint that countermands what the evidence calls for?