

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

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

3 DEDUCTIVE CONSTRAINTS: PROBLEM CASES, POSSIBLE SOLUTIONS

3.1 Intuitive Counterexamples

DEDUCTIVE consistency and deductive closure provide attractive constraints on ideally rational belief (for convenience, I'll combine these conditions under the heading "deductive cogency," or sometimes just "cogency"). The constraints of deductive cogency require, as we've seen, quite a specific conception of belief: a binary, yes-or-no attitude, which must consist in something over and above the agent's having a certain degree of confidence in the truth of the believed proposition. Presumably, if these constraints play an important role in epistemology, this role will be illuminated by an understanding of what the point of binary belief is. But before examining questions about the purpose or significance of this sort of belief, I'd like to look at some cases that directly challenge the legitimacy of taking rational belief to be subject to demands for deductive cogency. I think that the lessons these cases teach us prove useful in examining the question of whether the point of binary belief can motivate a cogency requirement.

Let us begin with a classic case often referred to as posing the "Preface Paradox."¹ We are to suppose that an apparently rational person has written a long non-fiction book—say, on history. The body of the book, as is typical, contains a large number of assertions. The author is highly confident in each of these assertions; moreover,

¹ A version of this argument was first advanced by Makinson (1965).

she has no hesitation in making them unqualifiedly, and would describe herself (and be described by others) as believing each of the book's many claims. But she knows enough about the difficulties of historical scholarship to realize that it is almost inevitable that at least a few of the claims she makes in the book are mistaken. She modestly acknowledges this in her preface, by saying that she believes that the book will be found to contain some errors, and she graciously invites those who discover the errors to set her straight.

The problem for deductive consistency is obvious. We naturally attribute to our author the belief, apparently expressed quite plainly in the preface, that the body of her book contains at least one error. We also naturally attribute to her beliefs in each of the propositions she asserts in the body of the book. Every one of these beliefs seems eminently rational. Yet the set of beliefs we have attributed to her is inconsistent. Moreover, the fact that our author, apparently quite reasonably, fails to believe that the body of her book is entirely error-free puts her in violation of the closure requirement.²

The problem here is clearly related to that posed by the lottery cases. There, if the agent believes of each ticket that it will lose, then he is precluded by consistency from believing that the lottery will have a winning ticket, and is required by closure to believe that it won't. But in at least one important way, the intuitive challenge posed by the preface case is sharper. In lottery cases, as we have seen, people do have some reluctance to assert flatly of their ticket that it will lose, and perhaps even to acknowledge believing that it will lose; this gives some encouragement to those who would deny belief—or rational belief—in these cases. But this is certainly not true of the individual claims made in the body of our author's book. Thus the dominant cogency-preserving response to preface cases

² In stating the preface case initially, I have been careful to be explicit about the fact that the belief expressed in the preface applies only to beliefs expressed in the body of the book, i.e. not to beliefs expressed in the preface itself. This is to avoid introducing complications of self-reference. In what follows, I will sometimes omit "the body of" for the sake of readability; I hope the intention is clear.

does not involve denying that the author rationally believes each of the claims in the body of the book.³

Defenders of cogency have thus typically wanted to deny that the author is rational in believing what I'll call the "Modest Preface Proposition":

Modest Preface Proposition. Errors will be found in the body of this book.

Denying rational belief in the Modest Preface Proposition clearly does not have the initial plausibility of denying rational belief in lottery-case propositions of the form "ticket *n* won't win." Admittedly, there would be something odd about a preface that baldly

³ An exception is Sharon Ryan's treatment of preface cases (1991), which argues that in all but certain very unusual cases, books by hard-working, intellectually responsible authors always contain unjustified claims. Ryan acknowledges that if one writes a short and simple book on addition for first graders one might succeed in writing a book with only justified claims in it; but in that case, of course, it does not seem intuitively that the modest preface statement would be rational. I doubt that this line can succeed in solving the preface problem for rational belief; it would seem to depend on setting the standards for rational belief excessively high. Given that responsible scholarship can easily produce rational beliefs about history (and not just about, e.g., elementary arithmetic), there is no barrier to producing history books consisting of rationally believed propositions. And given that rational belief need not be infallibly produced, a substantial book of such propositions may easily be highly likely to contain errors. One might object that, if we stick to Ryan's terminology of justified (rather than rational) belief, and interpret justification strongly, it is plausible that actual historians do typically make claims in their books that are not justified. Suppose this were granted. It still seems that cases posing the preface problem can be constructed easily. We might substitute for an academic historian a more humble sort of researcher: one who looks up telephone numbers, say, for a political campaign. I take it as uncontroversial that carefully looking up a person's phone number in the directory counts as a method of acquiring a justified belief as to what the person's phone number is. Now suppose that our campaign worker is incredibly scrupulous: she uses a ruler to line up the names and numbers, and she looks up each number on two separate occasions before entering it on her campaign list. It seems to me that, with respect to any particular number we might choose from her campaign list, she is justified in believing that it is correct. (If it is objected that memory limitations will preclude her from harboring hundreds of beliefs such as "Kelly Welly's number is 555-1717," we may concentrate on her beliefs such as "The 317th phone number on my list is correct.") Of course, phone directories are not infallible. Thus we need only make the campaign list long enough, and it will be overwhelmingly likely that there will be a mistake in it. And it seems that if our campaign worker understands this, she cannot be, as Ryan would have to claim, justified in believing her list to be error-free.

stated "This book contains errors!" But it does not seem at all odd to write in a preface, "In time, errors will be found in this book, and when they are, I hope that they will quickly be brought to light." And if an author is asked, "Do you believe that any errors will be found in your book?" or even, directly, "Will any errors be found in your book?" there is nothing at all unnatural about her saying, simply, "Yes."⁴

One might think that the problem could be avoided by taking a fairly lenient view of the demands made by formal constraints on rational belief. Taking deductive cogency as a rational ideal need not commit one to calling irrational anyone who falls short of the mark. One might, for instance, take the import of the cogency constraint to be something like this: If an agent can easily recognize that her beliefs are not cogent, and it is also clear how her beliefs could be revised to restore cogency, then rationality requires restoring cogency.⁵

In the preface case, however, the inconsistency is blatant. So is the lack of closure, insofar as we make the obvious supposition that our author lacks belief in what I'll call the Immodest Preface Proposition:

Immodest Preface Proposition. The body of this book is 100% error-free.

Moreover, it is clear that consistency can be restored simply by the author's dropping the belief that her book will be found to contain errors. And closure—insofar as the stipulated facts of the case go—could be accomplished by the author's adopting a belief in her own book's historical inerrancy. Thus, the violation of these constraints does not seem to be excusable, even on a moderate reading of the force of the constraints.

In addition, irrespective of how easily the departures from cogency could be discovered or repaired, the preface case does not

⁴ A quick check of the prefaces in books lying around my reading chair revealed the following sentences following immediately after an author's listing of those to whom he is philosophically indebted: "Their stimulus is largely responsible for what may be of interest in this book. The mistakes are all mine" (Mellor 1971, ix).

⁵ Maher advocates this sort of position on consistency (1993, 134–5).

seem to be the kind of case in which, even though there are certain improvements possible in the agent's beliefs, those improvements might rationally be forgone. The changes in the agent's beliefs here that would restore cogency do not strike us as possible improvements at all—they are as intuitively irrational as they are easy to formulate. Thus it seems that the preface case provides a strong *prima facie* argument against taking deductive cogency as a rational ideal, on any reading of how violations of formal constraints relate to rationality.

Finally, one might try to dissolve the difficulty that preface cases present by distinguishing carefully between first- and second-order beliefs. One might insist, for example, that closure would not require any second-order belief about the first-order beliefs expressed in the book. It might be conceded that closure would require belief in the conjunction of the book's first-order claims, but that belief might be held to be distinct, from the logical point of view, from the Immodest Preface Proposition.⁶

This line seems unpromising to me. For it seems clear that an author who knew what she had said in the body of her book could realize that this conjunction was materially equivalent to the second-order claim of inerrancy for the body of the book. Once she has accepted the equivalence, closure will take her from the conjunction to the second-order claim.⁷

One might try to block this line of reasoning by taking the normative force of formal principles of rationality to be conditioned

⁶ Simon Evnine (1999) uses something like this strategy in attempting to undermine a version of the preface problem framed as a challenge to the principle that rational beliefs are closed under conjunction.

⁷ Evnine concentrates on an extended version of the preface problem in which an agent reflects not just on the beliefs in a given book, but on all of her beliefs—"a vague, ill-defined and exceedingly large" set (1999, 206). This makes less transparent the relations between the first- and second-order beliefs. Evnine uses this to argue that lottery cases provide a stronger challenge to the closure-under-conjunction principle he ultimately seeks to defend. But given that book-oriented preface cases avoid the cited difficulty, and that lottery cases have their own difficulties (as noted above), it seems to me that the challenge presented by preface cases is stronger (indeed, Evnine's solution to the lottery problem does not apply to preface cases).

by the agent's limitations, so that closure would require that the agent believe propositions entailed by her other (rational) beliefs *only if the entailed propositions could be entertained by the agent*, in the sense that the agent could bring the content of the entailed proposition clearly before her mind. One might then argue that our author might well not be able to entertain the massive conjunction of all the claims in the body of her book, and thus that even if closure is taken as a rational ideal, our author is not required to believe in her book's inerrancy.

This objection also seems unpersuasive. It is undoubtedly true that ordinary humans cannot entertain book-length conjunctions. But surely, agents who do not share this fairly superficial limitation are easily conceived. And it seems just as wrong to say of such agents that they are rationally required to believe in the inerrancy of the books they write. Clearly, the reason that we think it would be wrong to require this sort of belief in ordinary humans has nothing to do with our limited capacity to entertain long conjunctions.

Moreover, even if we restrict the closure principle to entertainable propositions, restrict our attention to ordinary agents, and distinguish scrupulously between first- and second-order beliefs, the preface problem can be developed. Surely an ordinary author who was paying attention could entertain the conjunction of the first two claims in her book, and recognize the material equivalence of this conjunction and the claim

- (1) The first two claims in my book are true.

She would then be led by closure to believe (1). She could then easily entertain the conjunction of (1) and the third claim in her book. Our limited closure principle would then dictate believing that conjunction. Recognizing the equivalence of this believed conjunction with the claim

- (2) The first three claims in my book are true

would lead, by similar reasoning, to belief in (2), and so on, until the belief in her book's inerrancy is reached. It must be granted that only an agent hard-up for entertainment would embark on such a process. But it is certainly not beyond normal cognitive capabilities, and the inerrancy belief seems no less irrational for having been arrived at by such a laborious route.

3.2 Consistency without Closure?

Suppose it is granted that in preface cases it would be irrational for the author to believe that the body of her book is 100% error-free. Assuming that the author might yet be fully rational in believing each of the claims she makes in the body of her book, this would seem to require giving up closure. Still, it might be thought that ordinary ways of thinking and talking make the preface-based case against consistency somewhat weaker.

Although most authors would be highly reluctant to assert the inerrancy of their books—and not just out of false modesty—it is also true that many authors would be reluctant to assert “This book contains errors.” This might be taken as showing that authors typically lack belief in the Modest Preface Proposition. (It is, I think, unarguably natural to say “This book undoubtedly contains some errors.” But it might be claimed that “undoubtedly” signals that the agent is expressing a degree of confidence rather than binary belief. And it might be claimed that what explains some authors' reluctance to make the former statement is precisely that unqualified assertions express binary beliefs, and these authors lack the relevant belief.)

Of course, even if we accept the claim that the reluctant authors lack the relevant belief, this would not show that they were *rational* in withholding belief. But a position that mandated withholding belief in these cases might seem easier to swallow than one that required authors to have a positive belief in the inerrancy of their

own (current) scholarship. Might we salvage a partial defense of deductive constraints in the face of preface cases by retreating to the position that consistency is a rational requirement, even if closure is not?

This strategy avoids some of the implausible consequences of requiring full deductive cogency, but it seems to me that more than enough implausibility remains to undercut the value of the retreat. To see this, let us fill out a bit more fully the case of one particular “moderately immodest author”—one who does not assert (or believe) the Modest Preface Proposition, but who also does not assert (or believe) the Immodest Preface Proposition. This will allow us to see more clearly what the constraint of consistency by itself mandates in preface cases.

Let us suppose that Professor X, our moderately immodest author, sees himself as a solid historian. He would never write something in a book that he didn’t believe, or something for which he didn’t have very good evidence. But he also sees himself as a bit less neurotic than certain of his colleagues, in the following way: unlike them, he is free of a perverse fetish for endless minute and typically redundant fact-checking. He knows that each of his previous books has contained some minor errors of detail; this has, of course, allowed certain critics to exercise their nit-picking skills. But this does not bother Professor X much. After all, his books have been influential, and the broad conclusions they have reached are, he believes, entirely correct. Moreover, Professor X would point out that *every* book in the field—even those written by certain persnickety colleagues—has contained at least a few minor errors of detail. Indeed, he believes that writing a completely error-free book in his field is virtually impossible.

Given this background, it is not surprising that whenever a new book comes out—even a book written by a scholar he believes to be more meticulous than he himself is—Professor X believes that the new book will be found to contain errors. Time and time again, these beliefs have been borne out. And now, suppose that Professor X is studying a catalogue, in which his forthcoming book is being

advertised alongside the new offerings from Professors Y and Z (both of whom have taken unseemly pleasure in pointing out niggling little mistakes in Professor X’s previous works). He shows the catalogue to a nearby graduate student, chuckling, “I can’t wait until someone finds all the little mistakes in Y’s book.”

“You believe Y’s book has mistakes in it?”

“Of course I do. Why wouldn’t it?”

“Do you believe that mistakes will be found in Z’s book as well?”

“Yes! And I must admit, I’m looking forward to it. These anal-retentive types get so upset when they’re caught in the most trivial errors! Look—all my books have had some minor errors in them. But you see, that’s virtually inevitable, and it’s no big deal. I’m not as careful as Y or Z, but my reputation is—well, I’m sure you see what I mean . . .”

“So—your new book here—do you believe that it has any little errors in it?”

“No.”

Perhaps not everyone will share my intuitions here, but I think that Professor X’s last statement would strike most people as an obvious joke. And the reason for this is that to take this statement at face value would be to attribute gross irrationality to him. Given the comparisons our author willingly makes between his work and that of Professors Y and Z, it is simply not rational for him to believe their books to contain errors, but not to believe the same about his own book. The fact that withholding belief seems so clearly irrational here—that rationality would seem to require Professor X to believe the Modest Preface Proposition—provides powerful evidence that deductive consistency is not a rational requirement.

One might object that the last line in the above dialogue—a simple “no”—may be taken to indicate actual disbelief, rather than the weaker suspension of belief. It is true that “I don’t believe P” does often express belief in P’s negation, and not just an absence of positive belief that P. Thus, the defender of imposing consistency

but not closure might claim that the appearance of irrationality here is due to the impression that our author has adopted a belief in the Immodest Preface Proposition—a belief that is not required by consistency.

I think that this doubt can be dispelled by making the last line in the dialogue a bit more explicit. Suppose that, in answer to the student's asking if he believed that his new book contained any errors, Professor X had replied:

"No. I don't believe that my own book contains errors. I don't believe that it's error-free either. I'm just up in the air on that one."

It seems to me that this line is less funny only because it's pedantically drawn out. The attitude expressed by a literal reading of the dialogue is still absurd: our author believes quite firmly that every book in the field published so far, including his own, has contained multiple errors; he believes on these general grounds that Y's and Z's new books contain errors; he readily acknowledges that his own new book was written less carefully than Y's and Z's; and yet—somehow, unaccountably—when it comes to his own new book's inerrancy, he has no belief, one way or the other. This failure of his to draw the same conclusion about his own book that he so willingly draws about Y's and Z's books, when the evidence for the conclusion about his book seems to differ only in being somewhat stronger, strikes me as virtually a paradigm case of irrationality. Thus, I do not think that a defender of deductive consistency can escape the intuitive problem illustrated in the dialogue by noting that "I don't believe P" often means "I believe not-P."

Another objection to taking the dialogue as providing a serious intuitive challenge to consistency might be that its main character is too cavalier about getting things right to be ideally rational. Thus the fact that it seems wrong to impose consistency on *his* beliefs does not undermine the claim that consistency is required for ideal rationality.

I think that this objection misses the mark in two ways. First, it is not at all clear that our author's degree of caution in forming

historical beliefs falls short of the rational ideal. It does clearly fall below a Cartesian standard whose achievement would preclude even the slightest possibility of error. But the Cartesian standard is not the appropriate rational standard for historical beliefs. So, while it is clear that Professor X is not the most epistemically cautious person in his field, this does not show that his level of caution is sub-ideal.

Second, a similar situation could be constructed with a person who is at the epistemically cautious extreme in the field—say, Professor Y—as its central character. Professor Y might be more troubled by the inevitability of minor errors in history books. But she would presumably share certain key beliefs with her less cautious colleague: that all previous books in the field, including her own, had contained errors; and that all the new books by other very careful writers—Professor Z, for example—will be found to contain errors. If she refused to draw a parallel conclusion about her own new book, while acknowledging that she had no evidence that her current scholarship was more careful than her past scholarship (or that of Professor Z, etc.), then it would seem to me that this refusal to treat such epistemically similar cases on a par was clearly irrational.

Two final points should be mentioned in evaluating the strategy of responding to the preface problem by giving up closure and trying to save only consistency. First, in giving up on closure, one would lose a major part of the motivation cited by some defenders of imposing deductive constraints on binary belief. Pollock, for instance, takes his fundamental epistemological assumption to show that the epistemic importance of arguments requires a closure principle.⁸ Thus some defenders of deductive constraints would find the envisioned retreat unsatisfactory, even if it did avoid some sharply counterintuitive consequences.

Second, the motivation behind the retreat flowed from the intuitive strangeness of saying, flatly, "This book contains errors." But as

⁸ See Pollock (1983, 247 ff.); Kaplan (1996) argues for a similar point. This argument will be examined in detail in the next chapter.

mentioned above, it is not clear that unqualified assertions express belief, rather than claims to knowledge. Thus one might well want to explain some of the awkwardness of the flat assertion as flowing not from lack of belief, but from failure of a condition on knowledge that goes beyond rational belief.⁹

In sum, then, it seems to me that, while there may be something to be said for a position that imposes deductive consistency but not closure as a condition on rational belief, retreating to such a position does not help much to reduce the severity of the preface problem.¹⁰

3.3 Extent and Severity of the Intuitive Problem

It seems that the best response to the preface problem, if one wants to impose deductive constraints on belief, will involve biting the whole bullet: holding not only that our author should refrain from believing that his book contains errors, but that he should positively believe his book to be entirely error-free. Indeed, this is the tack taken by supporters of deductive constraints on binary belief such as Pollock, Maher, and Kaplan.¹¹ Before moving on to discuss how

⁹ See Williamson (1996, esp. sect. 3, and also 2000, ch. 11) for arguments that assertion is tied to knowledge rather than belief—even reasonable belief. DeRose's (1996) explanation of failure of assertability in lottery cases would seem to apply to preface cases at least as easily. On this view, our modest author realizes that if her book were, luckily, error-free she would have all the same evidence for her own fallibility, and thus would still believe that it contained errors. Thus she would judge herself not to *know* that her book contained errors, and would be unwilling to assert unqualifiedly that it did.

¹⁰ Another kind of retreat that would allow for intuitively rational beliefs in preface cases would be to impose only a limited consistency constraint. We saw that Kyburg (1970) showed that on a sub-certainty threshold view of belief, inconsistent sets of beliefs in lottery cases cannot be smaller than a certain size (where the size depends on the chosen threshold). Hawthorne and Bovens (1999, 241–64) make a similar point about preface cases. But as noted above, this sort of limited constraint is essentially just an artefact of the probabilistic constraints on degrees of belief; it does not provide an interesting independent principle for rational binary belief.

¹¹ In a recent treatment of preface cases, Adler (2002, ch. 7) explicitly claims only that the Modest Preface Proposition is not believed. But if I understand his position correctly, it would also sanction belief in the Immodest Preface Proposition.

one might either explain away the unintuitiveness or show that it must, in reflective equilibrium, be accepted, let us examine exactly what kind of unintuitive consequences the imposition of deductive cogency requires.

Consider a variant on the case examined above. Suppose that Professor X is a “fully immodest” author, who respects not only deductive consistency but deductive closure in the preface case. We'll join the dialogue part way through, after Professor X has expressed his firm beliefs that (1) every previous book in the field (including his own) has contained multiple errors; (2) he's not as careful a scholar as Y or Z; and (3) the new books by Professors Y and Z will be found to contain errors. Let's start at the point when the graduate student poses the crucial question:

“So—your new book here—do you believe that it has any little errors in it?”

“No. I believe that this book of mine is completely error-free.”

“Wow! Is that a first?”

“Absolutely. I believe that mine is the first book ever in the field that is 100% devoid of falsities.”

“Is this because your subject this time was particularly amenable to accurate scholarship?”

“Not at all.”

“Were you especially careful this time?”

“Certainly not. I'll leave the obsessing over trivia to Y and Z.”

“But doesn't all this make it pretty likely that there are at least some little mistakes?”

“Of course, it's overwhelmingly likely that my book contains many errors.”

“But you just said you believed . . .”

“Right. I believe that my book does not contain even one little error.”

I think that most people would be incredulous at Professor X's claims, if they took them as something other than a deliberate joke at his own expense. Taken literally, he attributes to himself a set of

beliefs that are, to my mind at least, patently irrational. The intuitive irrationality shows itself in at least two different ways.

(a) *Unequal Treatment*

The first way in which our author's beliefs seem intuitively irrational is a somewhat stronger version of the problem noted above in the first version of the example. Professor X comes to the conclusion, based solidly on excellent evidence, that errors will be found in other scholars' books. Simultaneously, he comes to exactly the opposite conclusion about his own book. Yet his evidence for errors in his own book seems even stronger than his evidence for errors in the books of others. Intuitively, this strikes us as irrationally treating similar cases differently. Given that the unequal treatment seems explicitly designed to favor the agent's epistemic assessment of his own beliefs, part of the intuitive irrationality here seems to spring from something akin to epistemic arrogance. But even if arrogance is not the agent's motivation, the unequal treatment seems indefensible.

One might object here that it isn't quite right to say that Professor X has better evidence for the existence of errors in his book than for errors in Y's and Z's books. After all, in the case of his own book, he already believes the claims it advances. Thus, these beliefs—which do entail the inerrancy of the book—give him a reason to think his own book to be error-free. And he clearly lacks a parallel reason for thinking the same of Y's and Z's books.

We should be careful not to allow this objection to sound stronger than it is. The objector cannot be claiming that Professor X has some reason for thinking that his own book is *less likely* to contain errors than are Y's and Z's books. In fact, as stipulated, our author is, quite rationally, more confident that his own book will contain errors than that Y's or Z's book will. So the unequal treatment is not a matter of our author's having any privileged reason for rational confidence in the truth of his book's claims.

To see this point clearly, consider what happens in cases where Professor X reads an authoritative book in a field somewhat distant

from his own. In many such cases, he straightforwardly believes what the book says. Let us consider such a case, in which Professor X's formation of new beliefs is a typical case of rationally accepting claims on authority. His acceptance of the book's authority does not, of course, mean that he thinks it to be an infallible source of truth. But it seems obvious that claims made in a book by a respected authority may meet the standards for rational belief.

Now, before reading the book, Professor X will, as before, have the reasonable belief that errors will be found in it. And after reading the whole book, his assessment of the probability of errors being found need not change at all. The book may not have an especially sloppy or especially careful style, and Professor X may have no special information that confirms or disconfirms the book's claims. Yet, the moment he reads (and believes) the final claim in the book, deductive cogency will require Professor X to execute an abrupt epistemic about-face, abandoning his original belief about the existence of errors in the book, and adopting instead the contrary belief that his colleague's book is 100% error-free!

The point here is not just to give another instance where deductive cogency demands intuitively irrational belief. The point is to make even clearer that the unequal treatment we saw in the original example cannot be justified by citing Professor X's special evidence for the beliefs in his own book. The reasons Professor X has for believing the claims he reads in this new book are no better than his reasons, before reading, for believing that the claims in the book were true. It's really just the bare fact that he now has adopted the beliefs expressed in the book that grounds his newly generous epistemic assessment of it.

Of course, there is an asymmetry between a case in which Professor X has read a new book, and one in which he hasn't yet read it. In the former case, once he has believed the book's claims, he does believe propositions that entail that the book is error-free. But that asymmetry cannot solve the intuitive problem facing the advocate of deductive cogency. For this asymmetry does not motivate differential beliefs about the existence of errors in the two books in any

way that is independent of the basic demand for deductive cogency. Perhaps there are independent arguments for deductive cogency that are strong enough to override its counterintuitive consequences. But whether or not this is true (an issue that will be examined in detail in the next chapter), it seems clear that the unequal treatment demanded by deductive cogency presents an acute intuitive difficulty.

(b) Internal Incongruity

The second way in which Professor X's beliefs in the example beginning the section exhibit intuitive irrationality is in the incongruity (reminiscent of Moore's Paradox) of "It's overwhelmingly likely that my book does contain many errors, but I believe that it doesn't contain even one."¹² Of course, it may well be that (as in Moore's cases) the sentence in question is not itself strictly inconsistent. But it is also clear that it is not a sentence one would expect to hear from any ordinary person—even a frankly arrogant scholar. In fact, it is hard to imagine anyone saying such a thing in ordinary life, at least with a straight face. (One can imagine it being said in jest: someone who has planned a 3 pm picnic and has just seen the morning weather forecast might say, "I know it's almost certain to keep raining all day, but I choose to believe that the sun will come out by 3 pm!" The humor here derives directly from the irrationality of the self-ascribed belief. If the person really did believe, in the face of all the evidence, that the sun would come out by 3 pm, we would not hesitate to deem her belief irrational.)

Moreover, not all of the incongruities are clear Moore-style examples mixing first-person belief-ascriptions with direct claims about the world. For Professor X's beliefs presumably will include both "My book very likely contains errors" and "My book does not

¹² A milder incongruity, even closer to Moore's Paradox, arises just from the imposition of consistency: "It's overwhelmingly likely that my book does contain many errors, but I don't believe that it does." Kaplan notes the Moore Paradox flavor of cases like the ones under discussion. His defense of the rationality of such beliefs will be examined in the next chapter.

contain errors." Like the more Moorean example, this does not quite constitute a contradiction; nevertheless, it certainly is not an intuitively rational combination of beliefs, and a person making both assertions categorically would strike anyone as bizarre.¹³

Thus it turns out that, in preface cases at least, imposing deductive cogency on rational belief conflicts quite dramatically with our ordinary practice. Ordinary rational people do not in such cases make the categorical assertions, or self-ascribe the beliefs, that deductive cogency would require. Nor would we be at all inclined to suspect that such beliefs were had by paradigmatically rational friends. In fact, the thought that the cogency-mandated beliefs are, or would be, rational in such cases is intuitively quite absurd. The bullet that must be bitten here is clearly substantial.¹⁴

3.4 Extent and Severity, cont.: Downstream Beliefs and Everyday Situations

Of course, if intuitive counterexamples to a general theory are few in number and peculiar in structure, we may have less reason to worry. Perhaps BonJour had this sort of point in mind when he consigned the preface problem to part of one footnote in a substantial book defending a coherence theory of justification which

¹³ This actually may understate the problem. As we will see in the next chapter, it is not clear how the two claims in question can rationally be believed by a cogent agent without giving rise to an explicit contradiction.

¹⁴ Some advocates of cogency are clearly concerned about this sort of intuitive problem, but not all. In Pollock's detailed treatments of preface cases (1986; 1990, ch. 9), he takes the problem to be just that of showing that the mechanism of "collective defeat"—by which he denies warranted belief in lottery cases—does not generalize in a way that would deny an author's warranted beliefs in the individual propositions asserted in her book. Pollock accomplishes this (in a way that stems from the fact that the lottery propositions are mutually negatively relevant, unlike the propositions asserted in a history book), and takes it to provide "a satisfactory resolution to the paradox of the preface" (1990, 253). Pollock does not seem to count it as part of the paradox that his account would bestow warrant on, e.g., our author's belief that his own book is the first 100% error-free contribution to the field.

includes deductive consistency as a necessary condition on coherence. After mentioning a different objection to imposing consistency, he writes: "And there are also worries such as the Preface Paradox. But while I think there might be something to be said for such views, the issues they raise are too complicated and remote to be entered into here" (1985, 240, fn. 7).

Yet it seems to me that the magnitude of the problem posed by preface cases should not be underestimated. One aspect of these cases that is not typically emphasized enough is that adopting a belief in the Immodest Preface Proposition is not something that occurs in an epistemic vacuum—especially for an agent who is deductively cogent. Let us think in concrete detail about Professor X's beliefs, to get a feel for some of the possible "downstream" effects of his believing the Immodest Preface Proposition.

In the situation envisaged, it would certainly be rational for Professor X to have the general belief that writing a completely error-free book would require being extremely careful and meticulous (which he knows he is not), or being amazingly lucky. Thus if Professor X's beliefs are to be deductively cogent, he must believe:

(a) I am amazingly lucky.

And believing himself amazingly lucky is only the tip of the iceberg. Given his belief that writing a completely error-free book would be such an unprecedented achievement, Professor X must take this fact into account in forming his beliefs about the future. For example, given his information about his colleagues, he is undoubtedly rational in believing that if no errors can be found in his book, Professors Y and Z will be in for a big surprise. In fact, he may well be rational in believing that if anyone wrote a completely error-free book in his field, it would soon lead to adulatory reviews, prestigious speaking engagements, and opportunities for professional advancement. He would then be committed to believing the following predictions:

(b) Professors Y and Z are in for a big surprise.

(c) I will soon receive adulatory reviews, prestigious speaking engagements, and opportunities for professional advancement.

Indeed, in many different situations, the Immodest Preface Proposition will combine with background beliefs to result in all manner of strange beliefs. Suppose, for example, we add to our story that, several decades ago, the Society for Historical Exactitude established a medal and a monetary prize, to be awarded to the first book in the field advancing substantial new theses in which no errors had been discovered one year after its publication. Although by now the monetary prize has grown to substantial proportions, the award has gone unclaimed, for reasons that Professor X understands only too well. Clearly, the fact that this award exists does nothing to weaken the requirement imposed by deductive cogency that Professor X believe the Immodest Preface Proposition. Now Professor X knows that his new book, like all of his works, advances substantial new theses. And since he believes that his new book is also the first such book to be 100% error-free, he should presumably believe:

(d) One year from now, I will be able to wear the SHE medal and start spending the SHE prize money.¹⁵

And this monetary belief will have consequences of its own. Given that Professor X knows himself to have particular automotive aspirations, he might well be justified in believing that if he comes into a sufficiently tidy sum, he'll buy an Alfa-Romeo. In that case, he should presumably now believe:

(e) In one year, I'll be driving an Alfa.

Clearly, it would be easy to pile on intuitively irrational beliefs in Professor X's case as long as one wanted to. And it's also clear

¹⁵ If the Society for Historical Exactitude seems too fanciful, it is worth noting that the connection between the feat of writing an error-free book in Professor X's field and the consequence of receiving large sums of money might easily be forged by more common mechanisms, such as job offers at well endowed universities. Thus, one could tell substantially the same story involving

(d') One year from now, I'll be working at a prestigious university and start spending my generous salary.

that it would require little ingenuity to come up with countless other cascades of intuitively irrational beliefs in different cases in which people obey deductive cogency by believing Immodest Preface Propositions. The structure of the problem involves cogency twice: first, cogency requires belief in an incredibly improbable proposition; then it requires belief in whatever propositions are entailed by conjoining the improbable proposition with ordinary reasonable background beliefs. The result is a chain reaction of cogency-mandated beliefs that are if anything more clearly irrational than the Immodest Preface Proposition itself.

Now it must be admitted that this plethora of intuitive irrationalities is almost certain to be short-lived. When the first reviews of his book appear, Professor X will come to realize that his book is not, after all, error-free; beliefs such as (a)–(e), which were spawned by the Immodest Preface Proposition, will then vanish as well. But, far from alleviating the intuitive strain that the example provides, this point should serve to underline an additional dimension of wackiness involved in holding beliefs such as (a)–(e). For a whole nest of beliefs like this will spring up anew each time Professor X publishes a book. Given deductive cogency, Professor X's knowing that similar sets of beliefs have arisen with the publication of each of his books, only to be leveled by the book's first reviews, will not serve in the slightest to diminish his epistemic duty to embrace each subsequent set of sanguine predictions.

Clearly, the problem posed by believing the Immodest Preface Proposition is not merely that this one belief is itself intuitively irrational. Adoption of this belief will have a strong tendency to spread ripples of intuitive irrationality throughout various parts of a deductively cogent agent's belief system. For the affected agent, the epistemic difficulties are quite severe.

It might be thought, though, that the malady, however grave, is at least a rare one. Might the defender of cogency take comfort in the claim that situations posing preface-like difficulties are rare, or remote in some way from ordinary epistemic life?

It must be admitted that actual prefaces worded in the way the classic example requires are fairly uncommon. But it is not at all uncommon, or remote from ordinary epistemic life, for people to write books expressing their beliefs. And many of these books are written in fields such as history or biography, where the number of factual details involved in a book makes it quite likely that the book contains errors. Whether or not these books have prefaces, deductive cogency would require all of their authors to believe them to be 100% error-free. Of course, very few of these authors have any such belief. And of those who have formed some belief on the question, the great majority undoubtedly believe that their books will be found to contain at least some minor errors. Thus, it turns out that there are quite a few real people in ordinary situations who have preface-style beliefs about their books—beliefs that, while intuitively quite rational, directly and obviously violate deductive cogency.

Moreover, beliefs of this structure are not restricted to authors of books in detail-oriented fields. Many of those who have reflected even briefly on their own fallibility believe that at least one of their (other) beliefs is mistaken.¹⁶ Some would undoubtedly hedge on the issue, saying only that they probably had at least one mistaken belief. But I suspect that only a tiny minority believe—as closure would dictate—that every single one of their beliefs on every topic is correct. Most would, I think, share Henry Kyburg's sentiment:

¹⁶ Again, the restriction to their "other" beliefs is intended to avoid self-reference. Evinne (1999) claims that this restriction, while needed to avoid self-referential paradox, poses a problem of its own. He claims that

*Something I believe, other than this belief, is false

"makes invidious distinctions among our beliefs and gives a special status to some that it does not give to others, namely, exemption from possible error" (p. 205). This objection should not, I think, be persuasive. After all, given that (as Evinne supposes) our reason for thinking ourselves fallible is empirical, it is not surprising that we have more reason to doubt some sorts of beliefs than others. People's beliefs about their hair color or addresses, for instance, are much less prone to error than their beliefs about, e.g., details in history. We in fact have excellent (empirical) reason to believe, about anyone, that she has at least one false first-order belief, and thus we have excellent (empirical) reason to believe that anyone who accepts the limited expression of epistemic modesty expressed by * is correct in so doing. So exempting * from its own scope is not, as Evinne claims, a case of "special pleading."

"I simply don't believe that everything I believe on good grounds is true, and I think it would be irrational for me to believe that" (Kyburg 1970, 59).¹⁷

Finally, apparently rational violations of deductive cogency may crop up even in quite ordinary situations that do not involve second-order beliefs. When I go to bed at night, I believe that the newspaper will be on my front porch by 6:30 the next morning. I don't, of course, have absolute certainty about this matter, but I've been taking the paper for years, and have more than enough experience of its reliability to make this a reasonable belief. I also have just as good reason to believe that the paper will be on my porch by 6:30 two mornings hence, and equally good reason to believe that it will be there three mornings hence, etc. If you ask me, "Do you believe that the paper will be on your porch by 6:30 am seventeen days from today?" I will answer affirmatively, without hesitation. I think it quite plausible to attribute to me, for each $n < 366$ at least, the belief that the paper will be on my porch by 6:30 am n days from now. But I also know that, on rare occasions, the paper does not arrive in the morning. Thus I also believe that, on some morning in the next year, the paper will fail to be there by 6:30. Clearly, these intuitively rational beliefs violate consistency.¹⁸

¹⁷ See also Klein (1985, 131); Kitcher (1992, 85); Foley (1993, 165); Nozick (1993, 78). Kyburg suggests that those who are tempted to deny the natural view here are misled by quantifier confusion: "of everything I believe, it is correct to say that I believe it to be true; but it is not correct to say that I believe everything I believe to be true." It is interesting to think about the downstream effects of following cogency rather than Kyburg here. If the implications of Professor X's belief about having written an error-free history book seem wild, they are surely tame compared with the consequences flowing from an agent's belief that she—and presumably she alone, among all the people who have ever lived—is correct about every single matter on which she has a belief.

¹⁸ This example is loosely based on an example in Hawthorne and Bovens (1999, 242). It might be objected that I do not really have the belief about, e.g., day 17 until I am asked, and thus, that if I haven't separately considered each day I do not have all the beliefs claimed for me in the example. Of course, we would not want to insist that for me to have a belief, I must be actively entertaining it. So the objection would have to be that I don't even have dispositional beliefs here. This objection seems weak to me, for two reasons. First, I think that we do typically attribute beliefs to people in propositions that they have not actually entertained, but which they would unhesitatingly agree with if asked. For example, I think that we would attribute to most people the belief that

Moreover, if I were to adopt the closure-mandated belief that, in this coming year, the paper would never once fail to be there by 6:30 am, my belief would be intuitively irrational.

The newspaper case may be thought to resemble lottery cases more than it does preface cases. But it is, I think, worth developing independently of the standard lottery case, for the following reason: in the standard lottery case, as we have seen, there is some intuitive reluctance to assert flatly "My ticket won't win," or to self-attribute the associated belief. In the newspaper case, my telling a house guest "The paper will be on the porch by 6:30 tomorrow" is entirely natural, as is my self-attribution of the correlated belief. Thus, our ordinary intuitive judgments about particular beliefs in the newspaper case seem to me to provide a clearer objection to deductive cogency than do our intuitions in the classic lottery case.

In sum, then, it seems that the intuitive challenge posed by apparently rational beliefs in preface-like and lottery-like situations is a strong one, in two dimensions. The beliefs demanded by cogency in some of these situations are not just slightly suspicious intuitively; they strike us as wildly irrational. And the situations in which intuitive rationality and deductive cogency conflict occur all the time, for ordinary agents in ordinary epistemic circumstances. Clearly, the intuitive burden imposed by deductive cogency cannot easily be shrugged off.

3.5 Undermining the Counterexamples?

Of course, even if it is acknowledged that our pre-theoretic judgments in the troublesome examples are firm ones, and that the counterintuitive cases are neither rare nor *recherché*, the existence of

there are more than 17 stars in the sky—even though not very many of them have had occurrent beliefs in this particular proposition. Second, I certainly could consider each of the day-specific newspaper propositions that figure in the example, and come to believe each, by any reasonable standard.

a large class of strikingly counterintuitive examples set in ordinary epistemic life does not in itself settle the issue. For one thing, it can always be argued that the reasons for imposing deductive cogency are so strong that our intuitive judgments about these examples should be overridden in reflective equilibrium. Assessing this sort of argument will depend on assessing the general arguments for imposing cogency; I'll turn to that in the next chapter. But a defender of deductive cogency might try a different tack. She might try to *undermine* (rather than override) our intuitive judgments in the apparent counterexamples, by showing, *on cogency-independent grounds*, why, e.g., it would be rational for Professor X to believe his book to be 100% error-free, or why it would be irrational for him to believe that some mistakes will be found in his book, or why I should not believe that the paper will be on my front porch tomorrow (or, perhaps, why I should believe that, this year at last, the paper delivery will never fail). Most such attempts I have seen focus on lottery examples; this is natural enough, since we do have some intuitive reluctance to claim belief of each ticket that it will not win. Below, I'll consider two such attempts, and will then examine one that focuses on preface cases.¹⁹

(a) *Guilt by Association*

One idea that has struck several authors as attractive in dealing with lottery cases focuses on the fact that such cases involve a set of beliefs (a) which, given what else is known (or rationally believed), contains at least one false member, and (b) whose member beliefs are very similarly based. The set of beliefs of the form "ticket *n* will lose" is of this sort. The idea is that, when the support that one has for each of a set of propositions does not significantly distinguish

¹⁹ There may be no very deep difference between arguments that seek to undermine our intuitive judgments in the troublesome cases on cogency-independent grounds, and those that seek to override our intuitive judgments in a way that depends on seeing cogency as essential to rational belief. I am separating them here mainly for expository convenience.

among them, and one believes that at least one of these propositions is false, then that support is insufficient for rational belief in any one of them.²⁰ Let's call this the Guilt by Association (GBA) principle.

Clearly, this principle will not help out with preface cases; it would be generally inapplicable (since the beliefs in the body of the book will not generally be supported in indistinguishable ways), and it would not yield the desired result anyway (we do not want to deny that authors can rationally believe the claims made in their books). So the GBA principle cannot be a complete answer to the intuitive problems with deductive cogency.²¹ Nevertheless, it does seem to have the advantage of meshing with our reluctance to attribute beliefs in lottery tickets' losing, and thus has some claim to providing non-question-begging motivation for preserving cogency in one important range of cases.

It is crucial to note that the principle is not a bare assertion of a no-known-inconsistency requirement; it does not come into play whenever one has a set of beliefs such that one knows that one of them is false. The beliefs in question have to be relevantly similar, so that, as Bonjour (1985, 236) puts it, the agent has "no relevant way of distinguishing" among the beliefs in the set.²²

²⁰ See Bonjour (1985, 235–6); Ryan (1996); and Evnine (1999). Ryan's version is a bit stronger, in that it forbids justified belief in the similarly supported propositions even in cases where one is not justified in believing that at least one of these propositions is false, but where one is justified in withholding belief about whether at least one of them is false. I should also note that, while Evnine is directly addressing rational belief per se, Bonjour and Ryan mean to be giving conditions on justified belief in the sense, roughly, of "meeting the justification condition of knowledge." I don't want to enter into the issue of whether this degree of justification corresponds to rational belief. Since I'm concentrating on the conditions for rational belief rather than knowledge, I'll just examine whether this basic approach can solve the problem that lottery cases pose for consistency constraints on rational belief.

²¹ It might generalize a bit beyond standard lottery cases. For example, it might be thought to help with the newspaper case, if one thought that the right reaction to that case involved denying that I am rational in believing that the paper will be on the front porch tomorrow.

²² A similar requirement is explicit in Evnine's statement (1999, 207). No such requirement is explicitly made in Ryan's official statement of her principle, but it seems implicit in her response to certain examples, and in her justificatory remarks on her principle (Ryan 1996, 130–5). Nelkin (2000) shows convincingly that Ryan's principle would be implausible without this requirement.

Now one worry one might have is expressed by Dana Nelkin (2000, 385): that principles of this sort are “so finely tailored to lottery-like cases that they are limited in their ability to explain what is really responsible for our lack of knowledge or rational belief in those cases.” Without some deeper motivation, the GBA approach might be dismissed by opponents of consistency requirements as an ad hoc response to lottery examples.

One might press the intuitive motivation question further by asking why it should be relevant that one’s grounds for the beliefs in question be very similar. It can’t be simply because the agent knows that the grounds can fail; after all, our whole problem arises only within a context in which we’re assuming that rational belief does not require infallible grounds. The thought must instead be something like this: “If an agent knows that a certain set of propositions contains a false member, she cannot rationally believe all of them. But since the grounds for believing the propositions are so similar, she has no non-arbitrary way of picking one not to believe. Thus, she cannot rationally believe any of them.” But if this is the motivation for the GBA strategy, it will not help at all to undermine the intuitive counterexample to deductive cogency. For it presupposes that it cannot be rational to believe a set of propositions when one knows that one of them is false. And this would seem to beg the question in favor of imposing deductive consistency.

Moreover, non-standard lottery cases reveal that the GBA approach does not even succeed at the limited task of squaring all lottery cases with consistency requirements. Consider a lottery in which different tickets have different (but always small) probabilities of winning. In such a lottery, there will be relevant differences among the propositions in the falsity-containing set. Thus the GBA principle will not apply, at least in a straightforward way.

Ryan (1996, 132–3) does consider an example of this general sort. In a million-ticket lottery where one knows that “the fifty blue tickets sold have a much higher probability of winning than all the rest,” Ryan holds that one must suspend judgment about the

blue tickets, but one may believe of the others that they won’t win.²³ But whether or not this is an intuitively reasonable thing to say in Ryan’s case, other unequal-probability cases will be harder for the GBA approach to handle. One might, after all, have a one-guaranteed-winner lottery in which each ticket had a different (but small) chance of winning. Here there is no set of competing beliefs that are epistemically indistinguishable. Thus the GBA principle would seem not to apply in this sort of lottery case at all.

If this is right, then it is even harder to see why we should think that the GBA approach provides a plausible way of defusing lottery-based counterexamples to deductive cogency. Principles that deny rational belief in lottery propositions may, as we’ve seen, derive some independent support from meshing with what we are intuitively inclined to say about lottery tickets. But it now seems unlikely that the GBA approach provides the correct explanation for, e.g., our reluctance to say “I believe my ticket will lose.” After all, we would be no less reluctant to say this sort of thing in an uneven-probability lottery case where the GBA principle is inapplicable. It seems, then, that those who want to undermine our intuitive judgments in the counterexamples to deductive cogency should look elsewhere.

(b) *Banning Purely Statistical Support*

Another lottery-inspired approach holds that a proposition may not be believed rationally if one’s grounds for belief are, in some sense, purely statistical. Of course, any consistency-saving approach to

²³ Ryan’s thought here, I take it, is that it is highly probable that some blue ticket will win. Thus, if one believed of each blue ticket that it would lose, one would have a set of relevantly similar beliefs which were such that one had good reason to believe that one of them was false. One cannot say the same of the white tickets, since it is not highly probable that one of them will win. It’s not clear to me that this way of handling the case fits with Ryan’s general objectives. She wants justified belief to serve as a sufficient condition (given truth and the absence of Gettier conditions) for knowledge. The present case does not seem like a Gettier case. Yet if the holder of a white ticket claims to know that her ticket will lose, and we find out that another white ticket wins, it does not seem that the knowledge claim was correct.

the lottery cases will have the consequence that high probability is not sufficient for rational belief. But some such accounts would, e.g., allow statistical support to suffice for rational belief in the absence of defeaters. The present idea is that statistical support is for some reason incapable of making belief in a proposition rational, even absent any special circumstances that might compromise that support in some way. Some writers (e.g. Cohen 1988, 106 ff.) have rejected statistically based beliefs as candidates for knowledge, and some have urged the same for rational belief. Here I'll concentrate on Nelkin's recent defense of the idea that statistical support is insufficient for rational belief.²⁴

One advantage of this sort of position over GBA is that it seems less ad hoc, in that it applies straightforwardly beyond the standard lottery cases. For example, it applies unproblematically to the unequal-probability lottery considered above. It also does not seem to beg the question; after all, it applies to cases that pose no threat to consistency requirements, such as lotteries in which there will probably be no winner. In many cases where we have statistical reasons for thinking a certain event to be highly improbable, we do seem reluctant to make flat assertions (or self-attribute beliefs) to the effect that the event will not occur. Insofar as this reluctance can be taken to show lack of rational belief, we have some independent motivation for the approach of banning purely statistical support (BPSS).

Of course, those who take assertability as tied to knowledge rather than belief will find this last motivation questionable.²⁵ And there is something at least curious in the basic BPSS idea. After all, no one thinks that statistical support is irrelevant to rational belief, and everyone acknowledges that it comes in degrees. Then why, one

²⁴ Nelkin (2000) clearly separates the belief and knowledge cases, and advocates related solutions for both. A related proposal for rational belief is made in Kaplan (1996, 125 ff.).

²⁵ DeRose (1996) notes the lack of assertability in lottery-like cases, but argues (1) that this is due to failure of apparent knowledge rather than failure of belief, and (2) that the failure of apparent knowledge is not due to the fact that support is statistical, but to violations of a counterfactual-based condition.

might wonder, can't it be sufficient support for rational belief in some cases?

Nelkin offers motivation for BPSS that goes beyond preserving deductive constraints or meshing with assertion practices. When an agent believes rationally, he can "see a causal or explanatory connection between his belief and the fact that makes it true" (2000, 396). Nelkin would not require the agent to give a detailed description of the causal/explanatory connection. But the agent must be able to posit the existence of such a connection: "the key idea is that ... I can take myself to believe something *because* it is true" (2000, 398). This seems to rule out rational belief that one's lottery ticket will lose, for example, because it is clear that the ticket's losing would not explain or cause one's belief.

Now one initial worry is suggested by the requirement that an agent "believe something *because* it is true." If Nelkin required a rational agent to believe that the fact that would make her belief true actually *caused* her belief, then many obviously rational beliefs would be deemed irrational. So, for example, having turned the flame on under my skillet three minutes ago, I now believe—without touching or otherwise examining the skillet further—that it is hot. But the skillet's hotness does not cause my belief (nor does it seem correct to say that I now believe the skillet to be hot because it is hot). My belief is causally or explanatorily connected with the fact that would make it true, but not in the simple sense of the fact causing or explaining my belief. In countless cases of rational belief, our belief that P is caused or explained by factors which, in turn, cause or explain the fact that P.²⁶

Let us interpret Nelkin's suggestion, then, to allow this indirect sort of causal/explanatory connection. On this interpretation,

²⁶ DeRose uses an example with a similar structure to reject Harman's claim that knowledge of P is made possible by an inference to the best explanation from one's evidence to P. DeRose points out that I can know, on the basis of reading in my copy of the paper that the Bulls won, that my neighbor's copy of the paper reports the Bulls' winning. DeRose points out that his own subjunctive account fares better here: if the neighbor's copy did not say that the Bulls won, then neither would mine, and so I would not believe as I actually do.

however, it is not clear that the suggested motivation for BPSS will succeed. Consider the agent's belief that his ticket will lose the lottery. This belief is explained by the agent's knowledge of the set-up and workings of the lottery in question. And it seems that this set-up was causally responsible for the lottery's outcome, including his ticket's losing. One might even take the fact that the lottery was set up in this way to provide an explanation of the fact that the agent's ticket did lose. After all, we do take the fact that a car is constructed in a certain way to explain the fact that it started when the ignition key was turned this morning (even though cars built this way do very occasionally fail to start). Of course, the issues surrounding the role of statistical connections in causation and explanation are extremely complex. But at this point, we have seen little reason to suppose that a causal/explanatory requirement will correctly weed out cases of statistically supported beliefs in a way that will help to motivate BPSS.

Moreover, it seems to me that BPSS faces a more acute problem that is independent of the motivational issue. Nelkin considers an example (from Harman 1986) in which Bill knows that Mary intends to be in New York tomorrow, and concludes from this that she will be in New York. But he also knows that Mary holds a lottery ticket with a one-in-a-million chance of winning, and that if her ticket wins, she'll go to Trenton instead. Harman's puzzle is about knowledge: Bill seems intuitively to know that Mary will be in New York, but not that Mary's ticket will lose. But the case provides, if anything, an even sharper challenge to Nelkin's view on rational belief. Intuitively, it seems quite rational for Bill to believe that Mary will be in New York. It also seems rational for Bill to believe that Mary will go to Trenton in the event that her ticket wins. Yet the BPSS approach does not allow us to say that Bill is rational to believe that Mary's ticket will lose, since that belief would have purely statistical support. Thus, the BPSS proponent is faced with the choice between embracing counterintuitive judgments or abandoning the deductive cogency requirement for rational belief.

Nelkin chooses to say that Bill lacks knowledge that Mary will be in New York. And although she discusses the case in terms of knowledge only, she would presumably deny that Bill's belief that Mary will be in New York is a rational one. But she hopes to mitigate the intuitive problem this sort of case poses by citing the rarity of such situations: "[I]t is important to note that we are not in Bill's situation very often. This means that it remains open that we often know where people will be (and not just where they are likely to be)" (Nelkin 2000, 407–8). Nelkin would presumably say the same about our rational beliefs about where people will be.

It seems to me, though, that denial of rational belief required here to maintain deductive cogency and BPSS will not be so easy to contain. It is true, of course, that we typically do not know that our friends are holding lottery tickets with the potential to derail their plans. But we do typically know, when a friend is driving to New York, that she'll be there only if she is not hit head-on by a drunk driver. Yet our grounds for believing that our friend won't be hit by a drunk driver seem to be of the purely statistical variety: we know that such events do sometimes occur, and we have no special reason to discount the possibility on our friend's particular route; however, we also know that accidents of this sort are incredibly infrequent. Moreover, it seems that in countless other cases our intuitively rational beliefs have this structure: we believe that *P* is true; we believe *P* is true only if *Q* is; and our reasons for believing that *Q* is true are merely statistical. We believe that we'll be at work on time; but we know we'll be late if (as occasionally but unpredictably happens) the bus breaks down. We believe that our car is parked where we left it; but we know that if (as occasionally but unpredictably happens) it has been stolen, it is somewhere else (see Vogel 1990 for this and several more examples). We believe that the Bulls won, but we know that if (as occasionally but unpredictably happens) the paper transposed the relevant scores, the other team has won (see DeRose 1996, 578–9). As Vogel points out, "Much of what we believe about the world beyond our immediate environments

could be made false by some chance event we haven't yet heard of" (1999, 166).²⁷

It turns out, then, that the BPSS strategy encounters a severe problem over and above the intuitive dubiousness of the idea that statistical support is somehow incapable of justifying rational beliefs. The strategy is particularly ill suited to providing a defense of deductive cogency. For it seems that, once one bans believing on purely statistical grounds, imposing deductive closure on rational belief forces one to embrace widespread skepticism—skepticism that vastly outruns any initial intuitive reluctance we have to claim belief in lottery-type propositions.²⁸ Yet it was just this intuitive reluctance that underlay the hope that BPSS could provide a way of independently motivating a cogency-respecting interpretation of lottery cases. Thus I think that the BPSS strategy cannot, after all, help to undermine these intuitive counterexamples to cogency.

(c) *Sorites, Commitment, and the Preface*

Preface cases present a harder problem for those who would undermine intuitive counterexamples to cogency. As we saw above, we do not attribute to ourselves the beliefs (and non-beliefs) required by cogency in preface cases; in fact, these beliefs (and non-beliefs) strike us as paradigmatically irrational. Thus, when faced with preface examples, defenders of cogency tend to argue at the

²⁷ Such examples are direct prima facie counterexamples to a closure principle on knowledge, since we are inclined to say that, e.g., we know who the President is, but don't know that he didn't have a fatal heart attack in the last five minutes. Vogel argues that this sort of example does not actually threaten closure; he argues that plausible ways of denying knowledge of the entailed proposition apply independently to the entailing proposition, so knowledge of the entailing proposition is not additionally threatened by denying knowledge of the statistically supported entailed proposition and then applying closure. For present purposes, there is no need to settle the issue of whether these examples involve some reason for disallowing rational belief other than that which would be provided by banning purely statistical support. If we have no special reason for disallowing beliefs supported on a "merely statistical" basis, then BPSS does not even get off the ground.

²⁸ See Vogel (1990), in which he terms this sort of threat "semi-skepticism."

general level that cogency must be a norm for rational belief, rather than trying to undermine our intuitive judgments in these specific cases. However, Simon Evnine (2001) provides cogency-independent arguments designed to show that it would be irrational for someone to believe that one of his (other) beliefs was false. In discussing Evnine's arguments, I'll adapt them to the specific case of the preface.

Evnine's first argument aims to show that the belief expressed in the problematic preface cannot be part of a fully rational set of beliefs. Let P be the Modest Preface Proposition: at least one of the claims in the body of this book is false; and let C_1 – C_n be the claims in the body of the book. Suppose (for reductio) that the author's beliefs in P and all of the C_i are rational. If each of the C_i is rationally believed, there can be no particular one of them in virtue of which belief in P is rational—in other words, none of the C_i individually is such that the author is rational to believe *it* false. But if that is true, then, Evnine claims, P should still be rational if we excise one of the claims—say, C_n —from the book. If we accept this, then we can repeat the steps of the argument, excising a belief at a time, until we get to the point where the body of the book is reduced to C_1 , and the author rationally believes P (which now applies, of course, to just C_1). But this is clearly absurd: it cannot be that both P and C_1 are rational; if the author is rational to believe C_1 , he's not rational to believe P .

Evnine gives several versions of this basic argument, including some designed for countably infinite belief sets. In each case, the argument assumes an analogue of the premise used in the version above: that subtracting a rational belief from the body of the book cannot make the preface belief irrational (or, conversely, that adding a rational belief cannot make the preface belief rational). But this sort of premise is clearly reminiscent of, e.g., the claim that plucking a hair cannot make one bald. The problem with such assumptions with respect to rational belief (rather than baldness) might best be demonstrated by considering analogous cases having nothing to do with deductive cogency. Arguments may be

constructed using essentially similar assumptions to derive radically skeptical conclusions about all sorts of intuitively rational beliefs.

Suppose, for example, that a murder takes place on a cruise ship. The ship's detective gathers all 317 people on board in the grand ballroom. The detective believes, on the basis of the extreme difficulty of anyone leaving the ship after the murder, that the murderer is in the room. It seems clear, on any non-skeptical view of rational belief, that the detective's belief may be rational. But the claim that this belief is rational would seem to be subject to a *reductio* exactly parallel to Evidine's *reductio* of the preface belief.

We may start by noting that there is no one person in virtue of whom the detective's belief is rational—i.e., there is no one particular person whom the detective rationally believes to be the murderer. So now, let us ask: can the detective simply excuse one of the 317 people from the ballroom, and remain rational in believing that the murderer is in the room? We may have some temptation to say “yes” to the first step here, if only because of the large number of suspects. But surely the general principle (“If the detective's belief is rational when n people are in the room, it will be rational when one of the n is excused”) must be rejected. For it would allow the detective to excuse passengers one by one until she rationally believed, of the last remaining passenger, that he was the murderer. Detective work just isn't this easy.

In many ordinary cases, an agent has a belief that at least one of a very large set of objects has a certain property, and the agent holds this belief on grounds that are not specific to any of the members of the set. And in any such case, one can offer a sorites-style slide into skepticism. If we accept such offers, we will end up admitting that we cannot rationally believe, e.g., that someone on the ship committed the murder; that at least one student at the University of Vermont was born in March; that at least one book in the library has a chapter that begins on page 17; that we ate spaghetti on at least one day in 1998; etc. Clearly, if there is something rationally defective

about the modest preface belief, it cannot be shown by this sort of argument.

Evidine also offers an independent way of explaining why we shouldn't believe the preface proposition. Unlike some other defenders of cogency, who admit that one must allow that it is probable that some of one's beliefs are false, Evidine denies even this. But then what about the inductive evidence provided by the beliefs of others, and one's own past beliefs? Evidine says that the inductive argument fails because one must be *committed* to one's own current beliefs, in a way that precludes thinking that some of them are false.²⁹

Aside from questions about whether we should see beliefs as commitments of any sort, it is unclear why the sort of commitment involved in belief would (or could) undermine the rationality of using inductive evidence in the ordinary way to support the Modest Preface Proposition. We are, after all, quite willing to form beliefs on less-than-conclusive grounds. The sort of commitment that would block even the moderate degree of epistemic modesty involved in believing the Modest Preface Proposition would seem appropriate only if our standards for rational belief-formation were Cartesian. Moreover, it is hard to see what, on this view, our attitude should be toward propositions related to the Modest Preface Proposition. One should not, presumably, believe that, unlike everyone else (and unlike one's former self), one currently is employing special methods of belief-formation of a uniquely reliable sort. Should one, then, simply be supremely confident that one is now astoundingly lucky (though one would

²⁹ Evidine compares believing to promising. One may have broken some of one's promises in the past, but, he asks, “Can one now address the promisee and say that one is confident that one will fail to keep some of the promises one is currently making?” (Evidine 2001, 167). The analogy here seems strained to me. In the promise case, the commitment involved makes sense in large part because it is the agent himself who makes it the case that his promise is kept. In the belief case, the agent clearly cannot make it the case that her beliefs turn out to be true. For this reason, the breaker of even a sincere promise is typically morally culpable for renegeing on his commitment. But the holder of a rational belief that turns out false has not thereby committed any epistemic sin.

appear to have no grounds for that assessment)? Or should one somehow refuse to form any opinion at all on how likely it is that one is possessed of special belief-forming methods or stunning epistemic luck? Nothing in the neighborhood of these thoughts seems even close to rational. And, more importantly, no such thoughts seem intuitively to flow from any sort of commitment one might undertake, merely in virtue of forming beliefs. Thus it seems to me that thinking about beliefs as involving commitments would do little to undermine our intuitive judgments in preface cases.

There are, no doubt, other ways of trying to undermine our intuitive judgments about rational belief in lottery and preface cases. But it is unlikely that defenders of cogency will succeed in removing the counterexamples' sting. Our pre-theoretic judgments, in the preface cases especially, are firm and stable. Thus, the best case for cogency will have to be made directly, and the consequent violence done to our pre-theoretic intuitions will have to be rationalized on the basis of the direct arguments for cogency. Let us, then, turn to examine those arguments.

4 ARGUMENTS FOR DEDUCTIVE COGENCY

It would be a mistake to dismiss deductive cogency merely on the basis of intuitive counterexamples, even if they are powerful and pervasive, and even if we see no way of undermining our intuitions in these cases. For it might turn out that anything we say on this topic will entail severe intuitive difficulties, and that rejecting cogency would carry an even greater cost than imposing it. After all, binary belief will, on any bifurcation account, be some propositional attitude whose point is not simply to reflect rational confidence in a proposition's truth. If we could be brought to see binary belief as an important and interesting component of epistemic rationality whose point requires deductive cogency, we might come to override our intuitions in the problematic cases.

In doing this, we might then seek some measure of reflective equilibrium by explaining the intuitions as resulting from a tendency to run binary belief together with another concept. For example, Mark Kaplan holds that our intuitive concept of belief is incoherent. We really have two separate notions: one of degree of confidence (or graded belief); and the other of acceptance (or binary belief). Graded belief alone figures in rational practical decisions. But binary belief has its own purposes, quite distinct from those of graded belief. And it is these purposes which lend importance to a variety of belief that is subject to the rational demand of deductive cogency.¹ Let us look, then, at some of

¹ Kaplan suggests that the Moore Paradox impression that one gets from assertions such as "I'm extremely confident that there are errors in my book, but I don't believe that there are any errors in it" stems from our confusing binary belief with a state of confidence. If belief in P doesn't require being confident that P is true, the sentence