

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

the arguments that have been offered in support of deductive cogency.

4.1 ... and Nothing but the Truth

A first stab at explaining why binary beliefs should be consistent flows from the very core of the concept of epistemic rationality: that epistemic rationality aims at accurate representation of the world. A natural expression of this idea as applied to binary belief is that an epistemically rational agent seeks to believe what is true, and to avoid believing what is false. If an agent's beliefs are inconsistent, she is automatically precluded from fully achieving the second of these objectives. How can this be ideally rational?²

One problem with this argument is that avoiding false belief is not the only epistemic desideratum: if it were, it would be rational to reject all beliefs. Having true beliefs is also important. Moreover, we do not even want avoidance of false beliefs to take lexical precedence over having true ones: if it did, it would be rational to believe only those propositions of whose truth we were absolutely certain. Yet once one sees rationality as involving a balance between the two desiderata, the quick argument for cogency collapses. For it would seem inevitable—on any weighting of the desiderata against each other—that there will arise situations in which the best balance between the desiderata will be achieved only by failing to maximize with respect to either one of them. In preface cases in particular, refusing to adopt the Modest Preface Belief (that mistakes will be found in one's book) keeps alive the logical possibility that one will avoid all false belief. But it is incredibly unlikely that,

isn't problematic; see Kaplan (1996, ch. 4). Maher (1993, 153) expresses a similar view about the folk concept of belief.

² Lehrer (1974, 203) makes this argument, though he's since given it up; see his (1975). Foley (1987, 257–8) has a very nice critical discussion of this sort of argument, partly along lines similar to some of those offered below.

in rejecting the Modest Preface Belief, one will avoid a false belief. The overwhelmingly likely consequence is that one will have forgone a true belief, and thus achieved a poorer balance of truth over falsity. Insofar as having true beliefs is desirable, the Modest Preface Belief looks like an excellent candidate for adoption.³

Moreover, even the goal of avoiding falsity—on any natural interpretation—itself militates against treating preface cases as cogency would require. Achieving cogency in preface cases requires the adoption of the Immodest Preface Belief (that one's book is 100% error-free). Now adopting this belief does, of course, leave open the logical possibility of perfect error avoidance. In fact, adopting this belief could not spoil an agent's perfect record of error avoidance—if she had one. But on any natural interpretation of the goal of error avoidance, it does not reduce merely to valuing error-free belief sets above others. It distinguishes among the other, imperfect, belief sets, and values having fewer errors (and, perhaps, less important ones) over having more. Since the Immodest Preface Belief is almost certain to be false, the goal of avoiding error will itself tell against this belief's adoption.

Still, the thought that the pursuit of truth will in some way rationalize deductive cogency may seem attractive. A sophisticated version of this type of thought seems in part to motivate Kaplan's assertion-based account of binary belief, which is designed to support a cogency requirement:

You count as believing *P* just if, were your sole aim to assert the truth (as it pertains to *P*), and your only options were to assert that *P*, assert that $\sim P$, or make neither assertion, you would prefer to assert that *P*. (Kaplan 1996, 109)

Given the considerations rehearsed above, it is not clear how this analysis would favor imposing cogency. Suppose, for example, we asked what the rational author of a history book would

³ Early examples of the basic decision-theoretic approach to binary belief are in Hempel (1960) and Levi (1967).

assert, with respect to the proposition that her book was error-free, if her sole aim were to assert the truth about this proposition. It would seem obvious—at first pass, anyway—that she should assert that her book was not error-free, since she is virtually certain that this proposition is true. But Kaplan does not intend the above-quoted passage to stand on its own; in particular, he wants to give a specific interpretation to the “aim to assert the truth”:

The truth is an error-free, comprehensive story of the world: for every hypothesis h , it either entails h or it entails $\sim h$ and it entails nothing false. This being so, the aim to assert the truth *tout court* is not one anyone can reasonably expect to achieve. But it is, nonetheless, an aim you can pursue—you can try to assert as comprehensive a part of that error-free story as you can. (Kaplan 1996, 111)

This interpretation of our epistemic goal—asserting as comprehensive a part of the error-free story as one can—draws our attention to the entire *body* of what one would be willing to assert, rather than to the individual propositions. Will this help us see how a cogency requirement could drop out of a desire to tell the whole truth and nothing but the truth? Would it license asserting a story containing the Immodest Preface statement, rather than an otherwise similar story containing the Modest Preface claim?

It seems to me that fixing our attention on whole stories does not in itself affect the argument significantly. True, refusing to assert the Modest Preface statement would leave open the bare logical possibility that the totality of one's assertions comprised a part of the error-free story. But of course, as Kaplan would certainly acknowledge, even this is not something anyone can reasonably expect to achieve. And if achieving total freedom from falsity is not a realistic option, it is hard to see what one would lose in foreclosing it. The total story that one ends up asserting is virtually certain to be a large part of the error-free story, plus a smaller budget of false claims. Including the Modest Preface Proposition in one's global story is virtually certain to increase the portion of the error-free

story one asserts, whereas including the Immodest Preface Proposition is virtually certain to increase one's budget of false assertions.

The idea, then, cannot be just to come as close as possible to telling the entire error-free story—at least, not in the obvious sense of maximizing the truth and minimizing the falsity in the story one does tell: there is no direct road here which starts from the desire to tell a story as close to the global truth as possible, and ends with the strictures of cogency. The idea must be that there is some independent sort of value in telling (or believing) a cogent story *per se*. Other authors have expressed something like this idea. Van Fraassen (1995, 349) writes: “The point of having beliefs is to construct a *single* (though in general incomplete) picture of what things are like.” Jonathan Roorda concurs:

our beliefs are not just isolated sentences in a collection; they are meant to hang together, to tell a univocal story about the way the world is. It is this feature of belief which subjects it to the requirement of deductive cogency: we do not require the gambler to make sure that all of the propositions he bets on be logically consistent; but we do require of the storyteller that the logical consequences of what she has already said will not be contradicted as the story unfolds.”⁴ (Roorda 1997, 148–149)

It seems worth emphasizing that, insofar as this sort of consideration is to support deductive cogency in a way that goes beyond the advice to believe only what is absolutely certain, we seem to have left the desire for accuracy behind in a fairly dramatic way. For it's not only that the value of telling a cogent story fails to *follow from* the value of telling the truth. Cases such as the preface show that defending cogency would require that the value of telling a cogent story actually *trump*, or *override*, the value of veracity. In the next section, we'll look at what might be said for this sort of view.

⁴ I should note each of these authors is defending an account of binary belief quite different from Kaplan's. On van Fraassen's account, belief entails certainty; on Roorda's, one believes only what one is certain of in at least some circumstances.

4.2 Keeping your Story Straight

Several writers have pointed out that we typically try to avoid asserting inconsistent sets of claims. An assertion-oriented account of belief, combined with the idea that rational belief is deductively cogent, would explain and justify this tendency. Now we've seen that the close ties between assertion and belief are contestable; assertions may represent knowledge claims rather than expressions of binary belief. But let's not pursue this issue here. Is there some special point in presenting to others (or even to one's self) a coherent picture of the world—a point that would outweigh the value of keeping the picture as accurate as possible?

It is important to keep in mind that our tendency toward maintaining deductive cogency in our assertions is far from absolute. Preface cases present dramatic examples in which our ordinary assertion practices violate cogency quite flagrantly. And if one ties assertion to belief, then our ordinary belief-attributing practices provide further violations of cogency; after all, everyday expressions of epistemic modesty such as “Everyone has some false beliefs, and I am not the sole exception” are made routinely, without upsetting ordinary believers/assertors in the slightest. Thus to use an assertion-based view to defend cogency, one would have to show that our ordinary beliefs (or assertion practices) in these cases are actually wrong for some reason. Is there something about the point of belief (or assertion) that makes our common-sense responses to these situations defective?

Kaplan, unlike many defenders of cogency, squarely addresses the intuitive challenge posed by preface cases. How, he asks, can the author of a history book be rational in asserting (and believing) that her book is 100% error-free, when she knows full well that this is extremely unlikely to be true?

In outline, the answer is quite simple. Unless she wants to give up the writing of history, our author has no choice... [O]ne simply cannot assert an ambitious, contentful piece of narrative and/or theory *without* running

a very high risk of asserting something false. So our historian has a choice. She can decide that the risk is intolerable, in which case she will refrain from writing history. Or she can decide to tolerate the risk and pursue her profession. (Kaplan 1996, 118)⁵

Now the point here cannot be that there is some real-world obstacle faced by professional historians who fail to assert, or to believe, Immodest Preface statements. (In fact, it's not clear that professional historians ever make the sort of Immodest Preface assertion required by cogency.) As Kaplan would surely acknowledge, the profession of history tolerates Modest Preface assertions without batting an eye. Still, it is clearly true, as Kaplan points out here, that one cannot typically assert a whole ambitious theory without asserting something that's likely to be false. And this could be turned into an argument for something like the claims quoted above. If we could show that asserting (or believing) whole ambitious theories was required for doing history (or systematic inquiry in general) correctly, then we could support the claim that Immodest Preface assertions (beliefs) were in some way intellectually necessary.

Is assertion of (or belief in) entire, highly detailed accounts of the world a necessary part of inquiry? Some doubt is engendered by the fact that actual Modest Preface statements vastly outnumber Immodest ones. One might try to explain away appearances here. It might be claimed that people are really thinking, inside, “My book is the first one ever to be 100% correct, even though my claims are no less controversial, and I haven't used especially reliable methods, and I haven't checked my facts more thoroughly, [etc.]...”—all the while disingenuously professing belief that the book will be found to contain errors, in order to appear modest. But

⁵ I should note that this is not Kaplan's main argument for cogency. One might even interpret the quoted claim as simply presupposing cogency: if cogency were mandatory, then there would be a sense in which the historian would be required to believe the Immodest Preface claim. But I think that Kaplan intends to be saying something more here. He follows up the quoted passage by arguing against certain alternative attitudes that inquirers might take to ambitious theories, suggesting that the quoted claim is intended to do more than point out that Immodest Preface assertions are required, given the presupposition of cogency. Kaplan's main argument for requiring cogency will be discussed in the next section.

this suggestion strikes me as psychologically implausible in the extreme.

It might also be urged that scientists do sometimes make unqualified assertions of large theories, or, perhaps more commonly, describe themselves as believers in large and detailed theories. But should we take these claims as assertions that the theories in question will never be found inaccurate in even the tiniest detail? Consider theories about the origins of the first human inhabitants of the Americas. Such theories surely do rest on extremely large collections of detailed claims about diverse matters, including linguistic data on current inhabitants of America and adjacent continents, genetic information, physical measurements and chemical dating of particular fossilized human remains, analysis of stone tools and other artefacts found in certain locations, dating of fossilized bones from animals apparently butchered by stone tools, claims about climatic conditions and animal extinctions (themselves based on various sorts of archeological evidence), and more. Now suppose an archeological anthropologist says that she believes a theory according to which the first Americans came from Siberia over the Bering land bridge. Will this support the view that rational scientists believe that the sort of large comprehensive theories under consideration are completely true?

If we take the "Siberian origins" theory to be just the single claim that the first American came from Siberia, then our scientist's belief is not of the sort under discussion, since her assessment of the probability of such a single claim may well be quite high. Her profession of belief will support Kaplan's position only if her theory is taken to include a large number of detailed claims, such as those involved in the interpretations of countless specific bits of evidence of the sorts mentioned above. In other words, the theory must include enough so that our scientist—like Kaplan's historian—will rationally be extremely confident that it is not completely true. Now suppose we make our question clear by asking her explicitly, "Do you believe that in the entire theory—including such-and-such details about this linguistic item being causally

connected with that one, this piece of rock being a tool fashioned at about such-and-such a time, the marks on this fossil being caused by scratches from a stone tool wielded by a human being, etc.—not one detail is incorrect?" Do we imagine that any scientist would answer affirmatively here? Or do we have any (pre-theoretic) intuition that it would be rational for her to have such a belief? If not, there seems little reason to think that rational scientists do really harbor the kind of beliefs that the argument in question sees as necessary for successful inquiry.

Moreover, in order to defend cogency as a rational requirement in the way envisaged, we would have to show more than that rational inquirers typically harbor beliefs to the effect that vast and richly detailed theories are completely true in every detail. It would have to be shown that this sort of belief plays a crucial role in inquiry. In other words, there would have to be some serious problems besetting inquirers who believed that even minor inaccuracies would ever be found in the details of their favorite ambitious theories. Scientists who merely believed that their favorite theories were *approximately* true in *most* respects would be at some sort of intellectual disadvantage. But it is hard to see what that disadvantage could be. Once a scientist or other inquirer has made all of the particular assertions involved in her ambitious theory, what is gained by her taking the extra step of asserting that her ambitious theory is absolutely flawless? It is hard to see what role such a performance would play in rational inquiry.

Of course, none of this is meant to deny that there is a scientific purpose in thinking and talking about big, detailed theories. There may well even be good reasons for scientists to form certain sorts of allegiances to such theories, to "work within" the systems that such theories provide. But all of this is quite compatible with believing that the theory one is developing is only approximately right: that the story it tells is largely correct; that the entities, processes, forces, events, etc., that it postulates are reasonably close to those that actually exist. In fact, it seems quite plausible to say that an important part of an inquirer's commitment to an ambitious theory

is precisely to identify and correct those parts of the theory that are mistaken!

The claim that acceptance of whole theories plays a crucial part in science has been made in a different way by Maher.⁶ He begins by endorsing Kuhn's observation that a highly successful large theory (paradigm) will not be rejected, even in the presence of anomalies, unless an alternative has been thought up. Maher notes that this may be explained by his account of acceptance: before development of the alternative, the anomalous evidence lowers the probability of the accepted theory. But since there is no alternative available, the fact that it offers a comprehensive account that's probably fairly close to the truth makes it rational to stick with it. When the alternative is dreamed up, there is a better option, and so the first theory is abandoned.

This seems entirely sensible, as far as it goes. But notice that the sort of commitment to theories invoked in this explanation need not include anything like the belief that the theory in question is true in every detail. The described commitment even seems compatible with the belief that the theory will be found to contain at least some inaccuracies. In fact, Maher notes that pre-Einsteinian physicists clinging to Newton's theory in the face of anomalies proposed *modifications* to Newton's theory, including modifications of the inverse-square law. This seems hardly the sort of behavior that would be expected of inquirers who believed Newton's theory correct in every detail, or even who refused to believe that it erred in any respect. Thus, while systematic inquiry may depend on investigators being guided by some sort of allegiance to a large theory, there is little reason to see this allegiance as incompatible with acknowledging the theory's imperfection.

⁶ See Maher (1993, 169 ff.). Maher's notion of acceptance, like Kaplan's, is supposed to capture an aspect of our folk notion of belief (the other aspect being degree of confidence). Maher takes this notion to be subject to a deductive consistency requirement.

So far, then, we have seen no reason to think that either our intuitive reluctance to assert Immodest Preface statements, or even our willingness to assert Modest Preface statements, is misguided. Nothing we have seen so far about the role of big theories in inquiry seems to give a point to our asserting or believing massively conjunctive claims which we rationally regard as highly improbable. If a mandate for imposing cogency on binary belief is to be supported by some part of our intellectual practice, we will have to find it elsewhere.

4.3 The Argument Argument

A third strand of argument intended to support cogency focuses directly on how logical relations seem to inform rational belief through *arguments*. John Pollock writes:

The main epistemological assumption from which everything else follows concerns the role of arguments in epistemological warrant. I have assumed that reasoning is a step-by-step process proceeding in terms of arguments and transmitting warrant from one step to the next in the argument. From this it follows that warrant is closed under deductive consequence ... (Pollock 1983, 247)

A "warranted" proposition, for Pollock, is one that an ideal reasoner would believe; he uses similar considerations to argue that warranted belief is deductively consistent. According to Pollock, arguments are as strong as their weakest links, and deductive inferences are completely warrant-preserving. Thus, a deductive argument from warranted premises must have a warranted conclusion.

Kaplan makes a similar claim about rational binary beliefs. He notes that, when a critic demonstrates via a *reductio* argument that the conjunction of an investigator's beliefs entails a contradiction,

the critic thereby demonstrates a defect in the investigator's set of beliefs—a defect so serious that it cannot be repaired except by the

investigator's abandonment of at least one of the beliefs on which the *reductio* relies.

But if it is a matter of indifference whether your set of beliefs satisfies Deductive Cogency, it is hard to see how *reductios* can possibly swing this sort of weight. (Kaplan 1996, 96)

Kaplan applies a parallel point to constructive arguments:

[I]f satisfying Deductive Cogency is of no moment, ... the fact that we convince someone of the truth of each of the premises of a valid argument would seem to provide her no reason whatsoever to believe its conclusion. (Kaplan 1996, 97)

Pollock takes arguments as justificatory structures within an agent's cognitive system. Kaplan puts his points in terms of the interpersonal persuasive force of arguments, but it's clear that he sees such force, when it is legitimate, as flowing from a rational demand on each of us to have deductively cogent beliefs. Thus, for both writers, the challenge of accounting for the rational force of *arguments* should be understood as the challenge of accounting for the way in which rational belief seems to be conditioned synchronically by deductive logic.

To evaluate this challenge, we should ask whether the rational force that arguments actually have can be explained *without* invoking a cogency-governed notion of binary belief. I'd like to begin examining this question by looking at something that may at first seem beside the point. Let us see how deductive logic constrains rational *degrees* of belief, in situations where we see arguments as doing serious justificatory work. As we saw in Chapter 2, there is a natural way of constraining rational graded beliefs that flows directly from the logical structures of, and relations among, propositions: it is to subject rational graded beliefs to a norm of probabilistic coherence.⁷

⁷ Of course, some would reject the notion of constraining graded beliefs in this way. For the present, I will assume (as do, e.g., Kaplan and Maher) that logical constraints on graded belief are legitimate—or, at least, that there is no in-principle objection to them that does not apply equally to cogency requirements on binary belief. In later chapters, I'll defend this assumption.

We can see right away that probabilistic coherence will force rational degrees of belief to respect *certain* deductive arguments: if P entails Q, then a rational agent's belief in Q must be at least as strong as her belief in P. Of course, this applies when P is a big conjunctive proposition. So if P is the conjunction of the premises of a valid deductive argument, and Q is its conclusion, then when a rational agent is very confident that the conjunction of the argument's premises P is true, she must believe the argument's conclusion Q at least as strongly. (If we think about this case in an interpersonal and diachronic way, we get a parallel result: my convincing someone to believe P strongly provides her with a reason to believe Q at least as strongly.)

The same mechanism works in *reductio* arguments. Consider a *reductio* aimed at rejecting Q, which is based on premises whose conjunction is P. In such a case, the conjunction P will be inconsistent with Q. Thus a rational agent's confidence in Q can be no higher than one-minus-her-confidence-in-P. So if she is confident to degree 9/10 in the conjunction of premises P, she must give no more than 1/10 degree of credence to Q. Similarly, her confidence in $\sim Q$ must be at least as high as her confidence in the conjunction of the *reductio's* premises. (Interpersonal-diachronically: if a critic points out to me that P is inconsistent with Q, and I am unwilling to give up my strong belief in P, I will have reason to give up my strong belief in Q.)

This sort of example shows that deductive arguments can have important effects on rational belief, even absent any cogency requirement—indeed, even absent any consideration of binary belief at all. So the Argument Argument does not come anywhere near showing that cogency requirements provide the only way for deductive arguments to gain epistemic purchase on us. But the examples above differ in two ways from what Pollock and Kaplan have in mind. First, they involve cases in which the agent not only finds each of the premises in an argument belief-worthy, but also finds the conjunction of the premises belief-worthy. Second, the belief-states described in the examples are graded, rather than

binary. Does the phenomenon that is evident in these examples extend to cover the sorts of cases that seem to prompt the pro-cogency argument?

Let us focus first on multiple-premise arguments. Suppose that there is a long argument from P_1 – P_n to C . In such a case, the fact that a rational agent believes each of the premises strongly does not necessarily give her any reason to believe C strongly. This is because having a high degree of confidence in the truth of each premise need not mean having a high degree of confidence in the conjunction of the premises. After all, the premises might be negatively relevant to one another, the truth of one making the truth of the others less likely (a simple example of this is seen in lottery cases, with claims of the form “ticket n won’t win”). And even if the premises are independent of one another, their conjunction will typically be far less likely to be true than any one of them. For example, suppose that P_1 is “the paper will be on my porch tomorrow morning,” P_2 is “the paper will be on my porch 2 mornings hence,” and so on. If we take the set of such propositions up to P_{365} , we get a valid argument for the conclusion “the paper will be on my porch every morning for the next year.” But probabilistic coherence does not force a rational agent who strongly believes each of the individual premises considered separately to believe the conclusion at all strongly. For believing each of the premises—even strongly—does not rationalize strongly believing their conjunction.⁸

Thus, it is clear that the logical force of deductive arguments on graded belief does not obey the principle that Pollock endorses: it is not generally the case that arguments are as strong as their weakest single links. In cases where one is not certain of the premises of an argument, we get the following result instead: a deductive argu-

⁸ The same point applies to *reductio* arguments. If $\{P_1$ – P_n , $C\}$ is an inconsistent set, this does not force one’s rational degree of belief in C to be low, unless one is not only highly confident in each of P_1 – P_n , but also confident in their conjunction. Since the issues below arise similarly for *reductio*s and constructive arguments, I won’t discuss *reductio*s separately in what follows.

ment that depends on a great many uncertain premises will (*ceteris paribus*) be significantly less powerful than an argument that depends on fewer. But this does not strike me as clashing with our ordinary ways of thinking about arguments. Surely we feel less compelled by an argument with a huge number of uncertain premises than by an argument with only a few—even if no particular one of the premises in the huge argument is, considered by itself, more dubious than the premises in the short argument.⁹

Of course, none of this shows that the effect of deductive arguments on degrees of belief exhausts the legitimate epistemic role of these arguments. And I suspect that those sympathetic to the Argument Argument will feel that the above discussion sidesteps the main issue completely. After all, what’s at issue is how deductive arguments affect rational *binary* belief, not how they affect rational

⁹ Pollock and Cruz (1999) present an example designed to challenge probabilistic analyses of arguments. They consider an engineer designing a bridge: “She will combine a vast amount of information about material strength, weather conditions, maximum load, costs of various construction techniques, and so forth, to compute the size a particular girder must be. These various bits of information are, presumably, independent of one another, so if the engineer combines 100 pieces of information, each with a probability of .99, the conjunction of that information has a probability of ... approximately .366. According to the probabilist, she would be precluded from using all of this information simultaneously in an inference—but then it would be impossible to build bridges” (p. 107). Here Pollock and Cruz seem to be endorsing the engineer’s simply relying on her conclusion about girder size to build the bridge—even though this conclusion is probably based on at least one false premise! At first, this might seem simply absurd, especially considering that the safety of the bridge depends on its girders having appropriate sizes. But of course, engineers do rely on the outputs of calculations with many inputs—inputs that are somewhat subject to error. Does this practice support something like Pollock’s “weakest link” principle? It seems to me that it does not. Engineers in the sort of situation envisaged presumably believe that any errors in their calculational inputs are highly unlikely to be large enough to affect the end result significantly. If this is right, then the real belief relied upon in bridge-building is not some (probably false) belief in the correctness of an exact size specification that follows from the multiple measurements. It is rather the belief that any errors in the input values are small enough that using the calculated value is close enough to be safe. (If the engineer did not believe this—e.g. if she believed that any errors in her premises were likely to be large enough to have a significant effect on her girder-size conclusion—then, if the probability of such an error even remotely approached the 0.634 level that the example specifies, using the calculated value in building a bridge would be unconscionably negligent.)

degrees of confidence. The arguments we have with others, which we write about in books such as this one, or rehearse to ourselves when we take a critical perspective on our own beliefs, are not overtly probabilistic. So the fact that deductive arguments can affect rational degrees of confidence might seem quite beside the point.

This protest would be decisive if rational binary belief were completely insensitive to rational degrees of confidence. But we have no reason to suppose that this should be so—in fact, quite the opposite is clearly the case. And insofar as rational binary beliefs are informed by rational degrees of confidence, the effects that deductive arguments have on the latter may well have important consequences for the former. This is particularly obvious on a threshold view, according to which binary beliefs just are graded beliefs of a certain strength. If one's confidence in the premises of an argument puts one's graded belief in the conclusion above the relevant threshold, it will thereby have produced exactly the effect we are looking for.¹⁰ Clearly, threshold views illustrate the possibility of deductive arguments affecting binary belief via their effects on graded belief.

Now threshold views are not the only binary belief model available; in fact, threshold models are often rejected explicitly by proponents of deductive cogency. Of course, in the present context it would beg the question to reject threshold accounts *because* they don't support cogency. But while cogency failure is surely the most common reason for rejecting threshold accounts of binary belief, these accounts may be criticized on cogency-independent grounds as well. For example, Kaplan (1996, 97–8) points out that if binary belief just is nothing more than a certain degree (call it n) of confidence, then it would be *impossible* (and not just irrational) to withhold belief in a proposition that one thought likely to degree greater than n . This would seem to render unintelligible (and not

¹⁰ Foley (1993, 167 ff.) makes some related points about *reductio*. He points out that the effectiveness of a *reductio* directed against one of a set of claims depends on both the size of the set and the strength of support for, and interdependence among, its members. See also Weintraub (2001).

just bad) Descartes' advice not to believe what is merely probable. For given one's degree of confidence in a proposition, the question of whether one believed it would already have been settled.

Kaplan's sort of worry could be answered by a metaphysically sophisticated first cousin of the threshold view. One might take binary belief, as Descartes apparently did, to be accomplished by an act of the will—an internal assenting to a proposition.¹¹ This would allow for the possibility of willing assent, or failing to will assent, to propositions in ways that did not match up with any particular level of probability. But one might further hold—not altogether implausibly—that *rational* assent (and thereby rational binary belief) was governed by a threshold of rational graded belief. (Indeed, this seems to be the form of Descartes' suggestion, with the threshold for rational belief set at absolute certainty.) This sort of view allows for the metaphysical bifurcation of binary and graded belief, while allowing deductive arguments to affect rational binary belief via their effects on rational graded belief.

So if the point of the Argument Argument is supposed to be that deductive reasoning can in principle play no role in conditioning rational binary belief unless binary belief is subject to cogency, then the argument is simply wrong. Our rational responses to deductive arguments may seem on the surface to flow from a cogency requirement; perhaps this helps explain why many epistemologists have seen cogency requirements as so obvious as to need no defense. But this interpretation of the role of arguments, initially appealing though it may be, is not the only one available. For arguments affect the degree of confidence it is rational to have in a proposition's truth; and, on virtually any account, rational degrees of confidence can have important implications for the rationality of binary belief. Thus, there is another clear route by which arguments may gain purchase on our rational binary beliefs. And this route is completely independent of any requirement of deductive cogency.

¹¹ See e.g. the Fourth Meditation (Cottingham *et al.* 1984, 37 ff.). This basic picture of belief is apparently older than Descartes; Derk Pereboom traces it back to the Stoics in his (1994).

Of course, the Argument Argument might be filled out in a less contentious way. It might be interpreted not as a claim that arguments can have *no* cogency-independent effect on rational belief. It might claim instead that the graded-belief-mediated effects of arguments are *insufficient* to explain the role that arguments legitimately play in our epistemic lives. Is it plausible that arguments legitimately affect us epistemically *only* via their effects on our graded beliefs?

Given the enormous variety of arguments, and of accounts of binary belief, it is hard to say much in general about this question.¹² And it would clearly be impossible to prove that no case exists in which a deductive argument rationally affects binary belief in a way that cannot be explained via the argument's effects on rational graded belief. But we can, I think, see that in some very typical instances, ordinary deductive arguments will have dramatic effects on rational graded beliefs, and these effects are just the sort we would expect in turn to affect binary beliefs—and affect them in exactly the way we traditionally associate with reasoning deductively. For example, suppose I look at my office answering machine, and form a very strong belief that

- (1) My office answering machine recorded a call as being from local number 865-4792 at 1:45.

I already have the following strong beliefs:

- (2) My house is the only one with local phone number 865-4792.
 (3) My wife, son, daughter, and I are the only ones who live at my house.
 (4) If 1 and 2, then someone called from my house at 1:45.

¹² We may generalize our point about threshold-style accounts a bit, and note that any account that allows degrees of rational confidence to provide a floor for rational binary belief, or a ceiling for rational non-belief, will be sensitive to arguments' effects on graded belief. And any account that does not do this must allow belief in vastly improbable propositions, or allow non-belief in virtual certainties. But the lottery and preface cases have already shown us that this price must be paid by any account of belief which protects deductive cogency requirements.

- (5) If 3, and someone called from my house at 1:45, then my wife, son, daughter, or I called from my house at 1:45.
 (6) I didn't call from my house at 1:45.
 (7) My son and daughter were at school at 1:45.
 (8) If 7, then neither my son nor my daughter called from my house at 1:45.

I then form the following strong belief:

C: My wife called from my house at 1:45.

Here we have a valid deductive argument with eight premises, each of which is necessary for deriving the conclusion. Moreover, each of the premises is something we would, in our ordinary binary belief-attributing practice, describe me as believing. And my belief in C is based on, and made rational by, my beliefs in 1–8. This seems to be a paradigm example of the sort of deductive reasoning we engage in daily. Can we account for examples like this by means of the argument's effect on rational graded belief?

Let us first think about how strongly I would rationally believe the premises. With respect to premises 1, 2, 3, 4, 6, and 8, I am virtually certain of each. The chance of my being wrong on any of these counts is surely less than 1 in 1,000; so let us set my degree of confidence in each of these, very conservatively, at 0.999. With respect to 5, there is some possibility that a call from my house to my office would have been placed by someone who didn't live at my house. True, I've received hundreds of calls from my house, none of which have come from anyone who didn't live there. But suppose we allow a very generous 1% chance of 5 being false, and set my degree of confidence in 5 at only 0.99. With respect to 7, it is possible that one of my children has, e.g., become sick at school, and has been brought home by my wife. But this certainly happens way less than once a year (and when it does, my wife lets me know as soon as she is called). Again, let us be very conservative, and set my degree of credence in 7 at only 0.99.

Now, as we've seen, the fact that each of these premises is itself highly probable does not entail that I must give high probability to their conjunction. But in the present case, there's no reason to think that the truth of any of the premises provides much reason for me to disbelieve any of the others. Let us suppose that they also don't lend one another significant mutual support, so that they are mutually independent. (This is of course not strictly true, but I think that they are independent enough so that we will not err too greatly in treating them as if they are. It is important to remember that we were extremely conservative in our original credence-assignments.) On this assumption, my credence in the conjunction of premises should be somewhat greater than 0.974. And it seems to me that this rough calculation passes the intuitive test: in the present sort of case, I should be very highly confident that all of 1-8 are true. Given this, of course, it follows immediately that my credence in C must be at least this high.

Of course, one example—even if it seems fairly typical of our day-to-day reasoning with deductive arguments—cannot refute the claim that there are other cases in which graded-belief effects cannot explain the legitimate power of deductive arguments. In general, cases in which rationally persuasive arguments can be understood as operating through graded beliefs are likely to be cases of arguments where we are very confident of the premises, where we don't have too many premises, and where the premises are positively relevant to one another—when they form an integrated, mutually supporting structure of claims—or are at least not negatively relevant to one another. In cases of these sorts, it will be rational to have reasonably high confidence in the conjunction of the premises of the argument, and, therefore, in the argument's conclusion.

In other cases, it is undeniably true that the graded-belief-based effect will be negligible. This will happen, for example, in arguments with large numbers of fairly uncertain or mutually negatively relevant premises. In these cases, results in conformity with the dictates of deductive cogency cannot be shown to flow from the

argument's effect on graded belief. But at least some of the clear cases in this category should give considerable pause to the advocate of the Argument Argument. For this category includes, paradigmatically, the very arguments where deductive cogency would lead us from reasonable premises to intuitively absurd conclusions: that my paper will never fail to be on my porch in the next year; that my history book is the very first error-free contribution to my field; that I, perhaps alone among all the people who have ever lived on Earth, believe only truths!

In response to the suggestion that the force of *reductios* depends on number of premises, Kaplan offers an example designed to show that even large *reductios* have rational force—force which, owing to the large number of premises, cannot be explained probabilistically.¹³ Kaplan asks us to suppose that he's been asked to produce a chronology of 26 events leading up to a serious accident. The chronology he produces has the following elements:

- P1: Event A preceded event B.
- P2: Event B preceded event C.
- P3: Event C preceded event D.
- .
- .
- .
- P26: Event Z preceded event A.

We point out that, given transitivity and non-reflexivity of temporal precedence (which he accepts), his chronology logically implies a contradiction. This *reductio*, Kaplan argues, “has critical bite: it exposes the fundamental inadequacy of the chronology I have produced.”

It is clear that the imagined chronology is fundamentally inadequate; but there are many sorts of inadequacy. If the example is to demonstrate the efficacy of certain *reductios*, we must, in assessing it, be clear about which claim is the target of the *reductio*. Perhaps the

¹³ Kaplan (2002: 459–60, fn. 20). Kaplan credits Ruth Weintraub and Scott Sturgeon for the suggestion to which he is responding.

most natural target to consider is the chronology itself, considered as one conjunctive claim. But this choice would demonstrate no problem at all; if the *reductio's* bite is to render belief in this great conjunction irrational, we have no difficulty explaining this bite on the basis of low probability.

In fact, similar points apply to other intuitively salient targets. Reconstructing important aspects of the events leading up to an accident will undoubtedly make use of transitivities to arrive at judgments about the temporal relations between, e.g., event D and event J. But if the events in which we're interested are the ends of a seven-link chain (as D and J are), the probability of the temporal-priority judgment we would naturally reach about them (that D preceded J) is less than $3/4$ ¹⁴. And if the events we're interested in are the ends of a 13-link chain, the chronology gives us no reason to place greater credence in either one being prior. Thus, when we think about many useful claims that might naturally be taken to be part of the chronology, there again seem to be good probabilistic explanations for the chronology's intuitive inadequacy.

This suggests that, insofar as the case poses a difficulty, it is with the individual elemental claims such as

P6: Event F happened before event G.

Indeed, these are the claims to whose probabilities Kaplan directs our attention: he points out that the situation described is compatible with our rationally having high (> 0.96) credence in each of these judgments. So perhaps the argument will be that although this level of confidence seems quite compatible with binary belief, the (large) *reductio* of P6 shows that P6 is not really belief-worthy. Since its unworthiness cannot be explained by the *reductio's* probability-lowering effect, cogency is required to account for the rational effect of argument here.

¹⁴ This assumes that we have no reason to trust some elements more than others, and that only one of the elements is false. If the events related by a particular judgment are the ends of a 7-link chain, then there's a $7/26$ chance that the error is in one of those links (in which case the judgment is false); thus, there is only a $19/26$ chance that it's true.

Now I think that it is not at all clear that P6 is unworthy of belief. But before thinking more carefully about this question, I'd like to fill the case out in a bit more detail, the better to fix our intuitions. First, the circular structure of the elements may be doing some intuitive work by suggesting that there is at least one *big* mistake in the elements. If that's right, we may well be influenced by the thought that the source from which we obtained evidence for the elements was not good, and that we shouldn't really be very confident of any of them. True, giving the elements 0.96 probability is mathematically possible in a case of this abstract structure. But that doesn't make 0.96 an intuitively realistic estimate of the probabilities in an actual case meeting the description. Let us, then, specify how the elements are arrived at. One might naturally imagine that the various events could somehow have been timed by relations to external events (e.g., the car was filled with gas at 8:15 am). But this would not lead to a circle of priority claims. To fix our intuitions as clearly as possible, let us try to fill in the abstract description in a fairly natural way, so that our evidence will lead to a circular structure with the high probabilities the argument requires. Here's one way of doing so (I've also taken the liberty of changing subject matter, to remove any distraction that might be caused if our supposition that we're reconstructing a serious accident for some important purpose had the effect of raising the intuitive standard for rational belief above 0.96).

Suppose there is a 26-person race, which we haven't seen. The rules stipulate that each racer will tell us who finished right behind her (and will tell us nothing else). The rules (which we may suppose are followed religiously) also stipulate that all racers tell us the truth, with the exception of the last finisher, who is to tell us that the racer who actually won finished behind her. We thus arrive at intuitively reasonable probabilities of 0.96+ for the elemental claims such as "racer C finished before racer D." Here, it is even less clear that belief in these claims would be irrational. Nevertheless, I think it must be acknowledged that many would be hesitant to assert unqualifiedly that D finished before E. And I think that some

would also deny the rationality of believing that D finished before E. (My own intuition, though not strong, is that this belief would be rational; but let us put that aside.) Supposing that we hold that such beliefs would not be rational, this could not be explained on probabilistic grounds. Would this show that there was, after all, a need to invoke cogency to explain why the beliefs weren't rational?

It seems to me that no such conclusion would be warranted. For once we have filled out the epistemic situation to rationalize the high probability judgments for the elemental beliefs, and once we have focused our attention on these elemental beliefs rather than on certain other beliefs that might be derived from them, the case very much resembles a standard lottery case. And this suggests that cogency demands may not be the best explanation of our reluctance to attribute rational beliefs. To test this suggestion, let us consider another race case, this one modified to remove the threatened *reductio*.

Suppose we receive 26 reports on who won each of 26 independent two-person races (e.g. "D beat K in race 6"). Suppose, however, that some of the people who report results are less than perfectly reliable—in fact, we know that, over a very long run, score-reporters have lied (with no discernible pattern) 1/26th of the time in this sort of context. Knowing this, what should we say if someone asks us who won race 6? I think that many would be reluctant simply to assert unqualifiedly that D beat K. And I think that those who were reluctant to countenance rational belief that racer D preceded racer E in the previous case would likely be just as reluctant to countenance rational belief that D beat K here.¹⁵

¹⁵ It is also worth remembering that, if the defender of cogency claims that one should believe the elements in the second race case, he must also hold that one should believe that, for some reason, the score-reporters told us the truth 26 times in a row! This is highly counterintuitive already (since on anyone's account, they probably haven't made 26 reports without lying); moreover, it is hard to see any reason for saying this in the present case that would not apply even if the reporters had given us 260 reports. Thus, insisting on rational belief in the second race case does not seem like an attractive option for the defender of cogency.

If that is correct, then it seems that the failure of rational belief in these two race cases should be explained in the same way. And the explanation does not flow in any obvious way from cogency, since in the second race case there is no *reductio*—no guarantee that one of the elements will be false. (This is related to the point that our reluctance in lottery cases to assert flatly, or avow belief in, claims such as "ticket 17 won't win" is not diminished when the lottery isn't guaranteed to have a winner.) Thus it seems to me that in the end we still have not seen a case in which the legitimate effect of argument on rational belief needs to be explained by a demand for deductive cogency.

Does this show that there are no examples that would serve the purpose of the Argument Argument? Certainly not. But those who would question deductive cogency requirements surely cannot be expected to demonstrate exhaustively that in every case where a deductive argument affects binary belief in an intuitively legitimate way, this effect can be explained independently of cogency. Surely the burden is on proponents of the Argument Argument to come up with specific, detailed examples of arguments whose rational efficacy cannot be explained in cogency-independent terms. For as we have seen, the general point that deductive arguments play a crucial epistemic role for us does not in itself establish a role for deductive cogency requirements.

It might be objected that I've underplayed the seriousness with which we actually take inconsistencies in the context of inquiry. Suppose, for example, the author of a history book were to discover that the claims *in the body of her book* formed an inconsistent set. Intuitively, wouldn't this be very disturbing? Might the fact that the crucial Modest Preface claim is, in some sense, "not really about history anyway"—that it oversteps, in some intuitive sense, the context of inquiry—make preface-type inconsistencies seem acceptable?

Now as we've seen, one can make a preface-like point with a great conjunction of the purely historical claims in the body of the book. And even bracketing this point, it is hard to see why an author

should be more concerned by an inconsistency within the body of the book than with preface-style inconsistency. After all, our comfort with the Modest Preface statement is directly based on our being highly confident that at least one of the claims in the body of the book is false. Discovering that the claims in the body of the book form an inconsistent set may elevate that high degree of confidence to certainty, but it is hard to see why this slight increase in our degree of confidence should be so alarming.

But wouldn't discovering inconsistency among the individual historical claims in the body of the book always *actually* be highly disturbing? I think that the answer to this question is less clear than it might seem at first. What the defender of cogency needs to make his point is a case involving an inconsistency that necessarily involves a great number of the huge and diverse set of historical claims making up the body of a book, and for my part I know of no case in which we've had experience of this sort of discovery in actual inquiry. Undoubtedly, people have found inconsistencies in the bodies of books, where the inconsistencies have been generated by a fairly small number of claims. But as we have seen, graded-belief-based effects may explain our felt need for revision in this sort of case. And in certain other cases, discovery of an inconsistency impugns one's general methods or sources in a way that significantly reduces one's confidence in some or all of the particular claims in the book. Again, however, our being disturbed in such cases can be explained in degree-of-belief terms. The kind of example that would bolster the argument for cogency would have to be one in which the discovery of the inconsistency did not significantly lower our confidence in the truth of any of the book's claims. Lacking experience with such cases, we cannot assume that they would actually strike us as calling urgently for epistemic repair.

Until persuasive specific examples are found, then, it seems to me that we've been given no good reason to think that deductive cogency requirements play an important part in epistemic rationality. Moreover, I think that at present we have at least some reason

for skepticism about the prospects for finding examples that will suit the Argument Argument's purposes. For any such example will have to be one in which we think that it is rational for someone to believe the conclusion of an argument based on the argument's premises, where all the premises are necessary to reach the conclusion, and yet where we also think that it's *not* rational for her to be confident that the premises are all true!¹⁶

Finally, it should be kept in mind that success for the Argument Argument would not be secured even by the discovery of a few cases that seem intuitively to fit this description. For the argument's success, there would have to be a considerable range of such cases. After all, aside from any intrinsic implausibility of the claim that rational beliefs may be based on premises that the agent is rationally quite confident are not all true, our intuitive verdicts on many cases provide powerful reasons to reject the demands of cogency. Thus, the cases adduced in support of the Argument Argument would have to be pervasive and persuasive enough to counterbalance the intuitive absurdities entailed by cogency requirements in the preface case, in newspaper-type cases, and in our ordinary expressions of epistemic humility.

In sum, then: there is certainly considerable surface plausibility to the idea that deductive arguments must derive their epistemic bite from deductive cogency requirements on binary belief. But it is also plain that submitting binary belief to cogency leaves us subject to bizarre arguments which run roughshod over our common-sense

¹⁶ This is not to beg the question by arguing that an intuitively persuasive example fitting this description should be disregarded just because it fits the description. It is intended merely to point out that examples of arguments whose effects cannot be accommodated by graded-belief-based mechanisms are going to resemble the strikingly counterintuitive applications of cogency in some respects—respects that are likely to make them counterintuitive as well. One might object that our intuitions in such cases would be distorted by our confusing binary belief with a state of confidence. But the objection would itself beg the question if it meant to argue that any intuitions based on rational graded belief must be discarded; after all, the degree to which rational binary belief depends on rational graded belief is part of what is at issue. If there are no cases in which an argument affects binary belief in a way that is very clearly correct intuitively, yet which cannot be explained via graded-belief effects, the Argument Argument is a non-starter.

understanding of rational belief. Insofar as there is an alternative way of grounding those deductive arguments which are intuitively legitimate contributors to epistemic rationality, we may resolve this tension nicely: we may maintain a healthy respect for rational argument without capitulating to the exorbitant demands of deductive cogency.

4-4 Rational Binary Belief

We customarily talk, and think, about our beliefs in binary terms. And it is certainly plausible to say that the point of beliefs is to represent the world accurately; that one's beliefs should comprise as much of the whole true story of the world as possible; that deductive arguments play an important role in determining which beliefs it is rational to have. But none of these observations about binary belief turn out to provide a sound motivation for a cogency requirement on binary belief.

Now this does not show that there is nothing to binary beliefs, or that there is no purpose to our talking about beliefs in an all-or-nothing way. It is clear that our everyday binary way of talking about beliefs has immense practical advantages over a system which insisted on some more fine-grained reporting of degrees of confidence. This is clear even if binary beliefs are understood on a simple threshold model.¹⁷ At a minimum, talking about people as believing, disbelieving, or withholding belief has at least as much point as do many of the imprecise ways we have of talking about things that can be described more precisely.

To take a trivial example, consider our practice of talking about dogs as big, small, and medium-sized. Obviously, talking about dogs in this way is extremely useful in everyday contexts. We would not

¹⁷ Foley, who defends a threshold model, makes a convincing case for the utility of binary belief-talk (Foley 1993, 170 ff.). Weintraub (2001) defends a threshold view along similar lines.

want to deny that, in a perfectly straightforward way, some dogs are big and some aren't, even if more precise ways of talking about dog sizes are available. And our rough sorting of dogs into three sizes even figures in everyday explanations: Andy provides good protection because he is big; Sassy is cheap to feed because she is small; etc. No one would advocate wholesale replacement of our everyday way of talking and thinking about dog sizes by some more precise metric—say, in terms of weights (or heights, or approximate weights or heights, or some function of approximate weights and heights). Any such wholesale change would clearly be counter-productive.

Nevertheless, as the example suggests, the obvious usefulness of talking about things using a given category doesn't show that the category "cuts nature at its joints." In the dog-size case, the interesting regularities—even the ones underlying the explanations mentioned above—will be more likely to be framed using more precise metrics. Small dogs do tend to eat less than big ones; but this regularity itself is explained by the way in which food consumption tends to increase with size, even within the "small dog" range. When we get serious about size-dependent effects—e.g. in calculating dosages of medicine—more precise metrics are quickly employed. Our rough-and-ready size categorizations do not seem to reflect the fundamental structure of the phenomena they describe.

Does our ordinary binary way of talking about beliefs pick out some epistemic property that's more important than bigness in dogs? Many epistemologists—even those, such as Foley, Maher, and Kaplan, who see graded beliefs as playing an important epistemic role—seem to think so. Kaplan, for example, considers a case in which you've just reported exactly how confident you are that a certain suspect committed a crime:

One of your colleagues turns to you and says, "I know you've already told us how confident you are that the lawyer did it. But tell us, do you *believe* she did it?" (Kaplan 1996, 89)

For Kaplan, there is something epistemically important left out when we give a description of a person's degrees of confidence.

For my own part, the colleague's question feels a lot like the question "I know you've told us that the dog weighs 79 pounds and is 21 inches high at the shoulder. But tell us: is it big?" When I "enter most intimately into what I call *myself*," I find no discrete inner accepting or "saying yes" to propositions. This seems particularly clear in cases where I move gradually from a state of low credence to a state of high credence (or vice-versa). I may start a picnic having heard a very positive weather forecast, and having no reservation about saying "I believe we're going to have a great picnic." But during the course of an hour, as clouds appear on the horizon and move toward us, as the sky gradually darkens, and as the breeze becomes stronger, my confidence in having a pleasant time fades, through the point where I no longer would self-ascribe the belief that we're going to have a great picnic, until, at the end of the hour, I would unhesitatingly say "I believe that our picnic is going to be spoiled." But at no point during the process do I seem to experience a discrete qualitative shift in my attitude toward the proposition that we'll have a great picnic—no jumps from an inner "saying yes" to an inner withholding of judgment to an inner "saying no." If, at some point in this process, I had said that I thought that the chances of our picnic being spoiled were 9 to 1, and someone asked, "But do you *believe* that our picnic will be spoiled?" I quite literally would not understand what information she was asking for.

Nevertheless, I don't mean to suggest that our binary belief talk is governed merely by degrees of confidence. As we've seen, we are somewhat reluctant to attribute beliefs in cases where the agent's high degree of confidence is based on blatantly statistical grounds. Various explanations of this fact may be offered: perhaps our belief-attribution practices are sensitive to some explanatory or tracking requirement, or to the fact that the statistical grounds somehow render salient the possibility of having the same evidence while being wrong. Adjudicating among these explanations would be a

substantial project in itself.¹⁸ But it is far from obvious that finding the right explanation—some rule or aspect of our belief-attributing practice that disqualified blatantly statistically based beliefs—would help reveal some binary state that was subject to interesting rational constraints (beyond those affecting degrees of confidence). This seems especially evident if our belief-attributing practice turns out to be sensitive to contextually determined conversational saliences.

To take another example, as Nozick (1993, 96–7) suggests, our willingness to attribute belief may depend on what practical matters are at stake. I would unhesitatingly describe myself as believing that our picnic will be a success if I were 98 percent confident that it would be a success and 2 percent confident that it would be spoiled by rain. But if I were only 98 percent confident that our airplane would arrive safely and 2 percent confident that it would crash, I would not unhesitatingly describe myself as believing we'd arrive safely. Again, various explanations of our practice are possible. For example, attribution of a belief that P might require the agent to have a high degree of confidence in P, but what counts as high might be sensitive to how badly wrong things could go if P is false. But again, without adjudicating among possible explanations, we can see how a practice of making black-and-white belief-reports that are sensitive to factors beyond degrees of confidence might make perfect sense, without its revealing any rationally interesting underlying epistemic state going beyond degree of confidence.

The project of working out the conditions under which people appropriately attribute binary beliefs may well reveal an interesting and complex pattern, even if our belief-attributing practice does not in the end correspond cleanly to a kind of state that is important from the point of view of epistemic rationality. After all, even the conditions under which we call dogs "big" may be interestingly

¹⁸ Writing about our willingness to make unqualified assertions, Kaplan comments: "Why we discriminate in these ways against matters of chance I am at a loss to say, but that we do seems quite clear" (1996, 127).

complex. Our practice there may not be governed by any fixed threshold of weight or height or weight/height combination. Perhaps factors such as the average size in some contextually relevant subset of all dogs, or the contextually specified use to which a particular dog is to be put, help determine our judgments. But working out complexities of this sort would not, I think, disclose any property that was important from the point of view of a systematic study of canine sizes.

The general reason for worrying that binary belief will not turn out to be an important part of epistemic rationality is this: insofar as our binary belief-attributing practices are sensitive to factors beyond rational degree of confidence in a proposition's truth, those practices are likely to point away from what we are most concerned with when we think about epistemic rationality. Let me illustrate with one clear example of this tension between going beyond rational degrees of confidence and maintaining epistemic importance. We saw earlier that Bonjour (1985), in response to standard lottery cases, denies that one is fully justified in believing that one's ticket will lose, no matter how high the probability is (though one may be fully justified in other beliefs whose probabilities are lower). But this move—exactly the sort needed by defenders of deductive cogency—seems to run directly counter, at least in spirit, to Bonjour's own characterization of epistemic justification:

[A]ny degree of epistemic justification, however small, must increase to a commensurate degree the chances that the belief in question is true (assuming that these are not already maximal), for otherwise it cannot qualify as epistemic justification at all. (Bonjour 1985, 8)

The worry is that there is no interesting notion of epistemic rationality that will sanction an agent's believing P but not sanction her believing Q, in a situation when she rationally believes that Q is more likely to be true than P.

This theoretical worry, of course, applies to virtually any bifurcated concept of binary belief. But if one also insists that rational binary beliefs be *deductively cogent*, then worries about the

significance of belief so understood multiply. Consider, for example, Professor X, our deductively cogent historian. We've seen how his belief in the Immodest Preface Proposition will commit him, given certain quite unexceptionable background beliefs, to believing some quite remarkable things: that he'll soon be receiving opportunities for professional advancement, that in one year he'll be enjoying a handsome salary and driving a brand-new Alfa-Romeo, etc. In Section 3.4, we saw that the intuitive irrationality of these beliefs (there labeled (a) through (e)) makes them *prima facie* counterexamples to cogency requirements. Here, I'd like to highlight a somewhat different angle: supposing that these *are* examples of rational binary beliefs, what do they reveal about the species of belief they exemplify?

Let us first think about how Professor X's beliefs should relate to the practical decisions he'll be making. Suppose, for instance, that he is offered an excellent deal on a new sensible car. His present sensible car could be nursed along for another year, so buying now will be quite disadvantageous if he buys a new Alfa-Romeo one year hence. On the other hand, if he does not receive the infusion of cash that would make the Alfa possible, he will do much better by taking advantage of the present offer. He believes, of course, that he'll be buying an Alfa in one year. Should he turn down the good deal on the sensible car? Obviously, he should not. The binary belief that he'll be buying the Alfa in one year, like various other beliefs that flow from the Immodest Preface belief, must be walled off carefully from Professor X's practical reasoning, lest he be led into countless idiotic practical decisions.

I should emphasize that defenders of deductive cogency requirements have sometimes said quite forthrightly that only graded belief should figure in practical deliberation. But this position seems much more palatable when one concentrates on just the Immodest Preface Proposition, whose obvious and direct practical implications are minimal. When the belief that one will be buying an Alfa in one year gets disconnected completely from the practical

question of whether to buy a car now, the point in having such a belief comes into question.

Moreover, disconnecting beliefs from practical reasoning in this way has bizarre implications for what one should *believe* about what one has practical reason to do. Suppose Professor X believes, as it seems he should, that

- (f) Anyone who has a perfectly decent car and is going to buy a new car in one year should not buy a new car now.

Given his beliefs about his own situation, deductive cogency would have him believing, quite rationally, that

- (g) I should not buy the new car now.

But this verdict must somehow cohere with the obvious fact about practical reasoning noted above: that Professor X would be quite irrational not to buy the new car now!

Moreover, the problem is not just that the deductively cogent agent's beliefs about his reasons for action are prised so far apart from what he actually has reason to do. A bit more exploration of these beliefs themselves raises serious doubts about the very coherence of cogency's demands. It seems obvious, for example, that Professor X should believe

- (h) It's very *unlikely* that I'll be able to afford an Alfa in the next few years.

But it's also hard to deny that he should believe

- (i) If it's very unlikely that I'll be able to afford an Alfa in the next few years, I should buy the new car now.

And given these beliefs, cogency would require Professor X to believe

- (j) I should buy the new car now.

The problem here is not, of course, that (j) is intuitively irrational—quite the opposite is true. The problem is that belief in

(j) is also *prohibited* by cogency, given Professor X's (cogency-mandated) belief in (g). It is not obvious just what beliefs a defender of cogency should recommend in this situation. Unless some way is found to deny the rationality of the beliefs leading to (g) or (j), it seems that cogency turns out to be unimplementable.

One way of avoiding this difficulty might be to argue that assertions and self-attributions of belief made in the context of practical decision-making did not express binary belief after all, but instead expressed degrees of confidence. I won't attempt to work out the intricacies of such an approach here. But I will note that, if there were a suitable way of circumscribing contexts of practical decision-making, it would effect a further corralling of the sort of belief to which cogency applied: such beliefs would end up being separated even from our ordinary ways of *thinking* and *talking* about practical decisions. (It's worth noting that this corralling of binary belief would have to exclude it even from certain contexts of inquiry. It's obvious that scientists, historians, etc., must make practical decisions in conducting their work—for example about expending research effort. But beliefs relevant to such decisions are subject to the sort of problem embodied above in Professor X's beliefs about whether he should buy the new car.)

In addition, there are many other ways in which the beliefs mandated by deductive cogency must be isolated from central parts of the agent's life. Let us ask: should Professor X be happy and excited that he will soon be enjoying a handsome salary and giving prestigious talks? Should he be surprised when he doesn't win the SHE prize? Presumably not. So, while it is obvious that one's emotions should in general be responsive to one's beliefs about the world, it is equally clear that they should not be responsive—at least not in any intuitively attractive way—to the sort of binary beliefs that would satisfy deductive cogency.

Now it might be claimed that disconnecting binary beliefs from emotive aspects of an agent's cognitive system is not much of an additional price to pay, once one has disconnected the beliefs from the agent's practical reasoning; after all, one might expect both

practical reason and emotions to be closely interconnected through the agent's values. But one might well think that if binary beliefs are to have any importance at all they cannot end up being a wheel that turns nothing in the cognitive life of the agent. What is left? As we have seen, one connection that is explicitly advocated by Kaplan is with the agent's assertive practice. Can deductively cogent binary belief find at least some of its importance in its connections with assertion (or, more specifically, with assertions motivated by the aim of truth-telling)?

It seems to me that even this connection is quite dubious. As we saw above, the rational assertability of the Immodest Preface statement was already intuitively suspect. It would be ridiculous (and, given how we ordinarily interpret assertions, even dishonest) for Professor X to assert in an NEH grant application that he had written the first error-free book in his field. And thinking about beliefs "downstream" from the Immodest Preface belief makes even clearer the strain that would accompany systematically connecting cogency-regulated belief to assertions. Consider assertions about the future. Should Professor X (insofar as he wants to tell the truth about the matter) assert to his friends that he'll soon be driving an Alfa-Romeo? Intuitively, it seems not—such an assertion would be highly misleading.

The connection is strained further when we consider what Professor X should assert about the practical decisions he or others need to make. Would it be reasonable for him to assert sincerely to a graduate student that trying to find even small errors in his book would be a waste of the student's time? Should he assert that he himself should not buy a new car now (even as he quite reasonably buys one)? Again, it does not seem that the kind of binary beliefs mandated by deductive cogency provide a basis for reasonable sincere assertion.

It turns out, then, that thinking about the sorts of binary beliefs required by cogency in some quite ordinary circumstances reveals strong reasons for doubting the importance of cogency-respecting binary belief. The reasons go beyond the general theoretical worry

about any bifurcated notion of belief: that such belief is determined by factors insufficiently related to truth. They concern the difficulty of connecting this sort of belief in any intuitive way with the rest of the agent's concerns, attitudes, or practices. To put the point another way, examples like those considered above raise the following sort of question: what point would there be in a practice of selecting a favored set of propositions to "believe," if this set of propositions included propositions of the sort Professor X is required to believe by deductive cogency?

Again, this is not to deny that our practice of binary belief-attribution is useful: clearly, it is. Moreover, it might well be interesting to see what governs this practice, and in so doing to gain insight into what we're talking about when we attribute binary belief. What is somewhat doubtful, though, is that this project will reveal to us a species of belief that will prove important from the point of view of epistemic rationality. And if we take binary belief to be a state that is governed by the constraints of deductive cogency, doubts about the state's epistemic importance become particularly acute. So while the ultimate interest of binary belief remains open to debate, it seems to me that if logic has a role to play in shaping epistemic rationality, it will not be the traditional one of subjecting binary belief to deductive cogency.