

Coming to Terms with our Human Fallibility: Christensen on the Preface

MARK KAPLAN

Indiana University

I. The Requirement of Deductive Cogency

Imagine that I have just finished delivering a philosophy paper. The question period has just begun. You raise your hand. Once called upon, you tell me that something I said very early on in my presentation, taken together with two things I said somewhat late in the presentation, logically imply a contradiction. What might I say in response? I might say that I actually hadn't said all these three things. Or, I might deny that the three things logically imply a contradiction. What I *wouldn't* do—because it would be absolutely outrageous for me to do it—is admit that I said the three things, admit that they logically imply a contradiction, but deny that this was any criticism of what I said. Thus the following principle:

(The Requirement of Consistency) The set of things you are willing to say in the context of inquiry should be (in the sense that you open your position to criticism if the set is not) consistent.

But, actually, a stronger principle is warranted. Suppose that you did something different: you pointed to these three things I've said, and observed that they jointly logically imply P. And then you challenged the propriety of my being willing, perhaps of anybody's being willing, to say that P. No charge of inconsistency being leveled here, yet a legitimate form of challenge all the same. I might try to meet the challenge by arguing that the very things to which you've pointed provide a good reason for saying that P. Or, I might seek to meet your challenge to the propriety of saying that P on independent grounds. But what I *wouldn't* do—because it would be outrageous for me to do it—is say,

“Oh, you’ve misunderstood me. I’ve said the things you’ve pointed to. And, true, taken together, they logically imply P. But I’m not *saying* that P. And, only if I were *saying* that P, would I have any call to meet your challenge to the propriety of being willing to say that P.” In saying what I originally did by way of delivering my paper, I committed myself (though, perhaps, up to this point unwittingly) to saying that P. And so my saying what I originally did is open to legitimate criticism if (for example) there is an excellent reason why one should not want to say that P. Thus a stronger principle:

(The Requirement of Deductive Cogency) The set of things you are willing to say in the context of inquiry should be (in the sense that you open your position to criticism if the set is not) consistent, and contain (at the very least) every logical consequence that you acknowledge to be (i.e., that is, and that you are willing in the context of inquiry to say is) a logical consequence of what the set contains.

It is important to notice that the Requirement of Deductive Cogency concerns itself solely with the conditions under which you open your position—that *this* is the set of things you are willing to say in the context of inquiry—to criticism. The requirement says nothing about when *you* are open to criticism for its being your position. The conditions whose satisfaction is necessary to satisfy the requirement are far too demanding for it to be thought that your position’s violation of those conditions automatically opens *you* to criticism for its being your position.¹ But, for all that, the requirement still has bite. After all, the question as to what renders a position open to criticism is central to the enterprise of rational inquiry—an enterprise whose aim is, in part, to determine what positions may be sufficiently immune to criticism to warrant adoption. The Requirement of Deductive Cogency is meant to offer a partial answer to that question.²

¹ You can hardly be held open to criticism for (say) being willing to say an inconsistent set of things whose inconsistency you are unable (and most in your place would be unable) to see. All the same, the set *is* inconsistent and so your position is open to criticism on that account. Nor could you plausibly be held open to criticism if, upon realizing that the set is inconsistent, you continued to hold (if somewhat gingerly) onto your position until you had a decent opportunity to see how that position might best be modified to remove its vulnerability to *reductio*—and to see how costly (in terms of the epistemic benefits the position affords) such a modification might be. Indeed, it may well be that, if the cost were high enough, the reasonable thing for you to do would be to maintain the position (being careful, of course, not to exploit its inconsistency in untoward ways) until a less costly modification became available. The Requirement of Deductive Cogency leaves that possibility open.

² For more on this, see Kaplan 1996, pp. 36–8, 113–4.

The trouble is that the Requirement of Deductive Cogency would seem to be decisively undermined by an obvious and devastating objection: the Requirement of Deductive Cogency holds that a person opens her position to criticism unless, upon completing a long, ambitious book of historical scholarship, she is prepared to say that everything she has said in the book is true. The thought is that it would be mad for her to be willing to say something so improbable, so outrageously immodest.

Now, it could be argued that this actually is no objection at all to the Requirement of Deductive Cogency. After all, the Requirement of Deductive Cogency doesn't say that the right thing for our author to do, all things considered, is to be willing to say that everything she's said her book is true. All the requirement says is that she opens her position to criticism if she isn't willing to say that everything she's said in the book is true. It is entirely compatible with this that she also opens her position to criticism (the criticism that what she's willing to say is so very improbable and immodest) if she *is* willing to say that everything she's said in the book is true. In that event, determining what is the right thing for our author to do, all things considered, would come down to determining which of the two criticisms is the more serious. The verdict that it would be mad for our author to be willing to say that everything she's said in her book is true, it could be argued, is to be understood as simply reflecting the judgment that the second criticism is much more serious than the first.

But this line of defense won't do. It's not because the verdict cannot be understood this way. It is, rather, because the verdict can just as easily be understood as reflecting the judgment that our author opens her position to *no criticism at all* if she is unwilling to say that everything she's said in her book is true. Now, to be sure, we have, in favor of understanding the verdict in the former way, the consideration I offered earlier in support of the Requirement of Deductive Cogency—that it would be outrageous for me to treat as no criticism at all someone's demonstrating (and to my satisfaction) that three things I have said in a presentation logically imply something I am unwilling to say. But fair is fair. If it would be mad for our author to be willing to say that everything she's said in her book is true, then that consideration should count just as strongly in favor of understanding the verdict in the latter way. The moral would seem to be that a defender of the Requirement of Deductive Cogency simply cannot afford an indifference to whether it would be mad of our author to be willing to say that everything she's said in her book true.

So I want to explore a more direct line of defense. I want to examine whether it really is true that it would be mad for our author to be

willing to say that everything she's said in her book is true. In particular, I want to examine David Christensen's attempt, spread over some three chapters of his elegant little book, *Putting Logic in its Place: Formal Constraints on Rational Belief*,³ to convince us that it really is true. My narrow purpose in so doing is to mount a defense of the Requirement of Deductive Cogency. But I also have a broader purpose.

Christensen's major argument for the Requirement of Deductive Cogency's having mad consequences is, I want to argue, an argument for a form of skepticism. By this I do not mean the skepticism Christensen means to express: skepticism as to the propriety of the Requirement of Deductive Cogency. Rather, I mean that Christensen's major argument is, in effect, an argument for a skepticism of the traditional philosophical sort: a skepticism that says, of some epistemic entitlement we act as if we have, that this is an entitlement that we do not, and could not possibly, have. My larger purpose is to explain why his argument constitutes an argument for skepticism of this sort—it is, I will contend, a skeptical argument from the fact of our human fallibility—and how the argument goes wrong. Hence the promise implicit in the title of this essay: that I mean to say something substantive about how, as philosophers, we are to come to terms with our human fallibility.

II. The Confidence Threshold Requirement

But before I do all this, I want to be as clear as I can about why I think we simply cannot make sense of our intellectual lives unless we grant the propriety of the Requirement of Deductive Cogency. I want to do this because I know there are those who will think that there is available a straightforward, simple and perfectly adequate account of what is going on, when we decide what we are willing to say in the context of inquiry, that rejects the Requirement of Deductive Cogency and avoids entirely its *prima facie* mad consequence.

According to this account, deciding what you should be willing to say in the context of inquiry comes to nothing more than deciding what you are warranted in believing. And what is to believe P? It is just to be suitably confident that P, where that requires being more confident than not that P, but does not require being maximally confident that P—that is to say, as confident as presumably you are in (P or not-P). So deciding whether you should be willing in the context of inquiry to say that P comes to nothing more than deciding whether you are

³ Christensen, 2004, chs 2–4. All parenthetical page references that appear in the text from this point on are references to this work.

warranted in having sufficient confidence in the truth of P as to exceed the threshold that must be exceeded for you to count as believing P.⁴

Now, according to this account, there is an important constraint that your investments of confidence in propositions have to satisfy if they are to count as warranted: they have to answer in substantial way to the axioms of the probability calculus. That is to say, it is a condition on your investments of confidence in propositions being warranted that, at the very least, the propositions upon which your investments succeed in imposing a rank (according to which you are more confident are true, and which less) could be so ranked in probability. Thus, according to this account, it is a condition of adequacy, on any principle that would impose a constraint on what you should be willing to say in the context of inquiry, that this constraint be a legitimate constraint to impose on the propositions in which you are warranted in investing confidence above the threshold for belief, where your investments of confidence answer in the foregoing sense to the axioms of probability.

But this is, of course, a condition of adequacy that neither of the requirements I put forward satisfies. As the lottery paradox reminds us, you can be perfectly warranted in investing as much confidence as you like short of the maximum in each member of an inconsistent set of propositions. So, according to this account, both the Requirement of Consistency and the Requirement of Deductive Cogency must be rejected.

Now, I do not want to deny that the account just described has its appeal. I think it is pretty clear that we find it natural to think of belief as a state of high confidence. And I think there is compelling reason to agree that it is condition on our investments of confidence in propositions being warranted that these investments satisfy the aforementioned probabilistic constraint.⁵ I also think that there is no question but that we talk as if the matter of determining what we are warranted in being willing to say in the context of inquiry comes to nothing more than determining what we are warranted in believing. Witness the fact that, in the context of inquiry, we treat as entirely interchangeable talk of whether we believe that P, and talk of whether we are willing to say (or are saying, or want to say) that P. Witness, as well, the fact that we find it natural to suppose, in posing the preface paradox that, once we have learned how the things we are willing to say in the body of a

⁴ I leave open the matter of whether the confidence threshold might vary from proposition to proposition and/or from context to context. So long as exceeding the threshold requires being more confident than not, but does not require being as confident as one is in P or not-P, the matter of whether (and, if so, how) the confidence threshold might vary is of no import to what follows.

⁵ I explain why in the first chapter of Kaplan 1996.

scholarly book constrain what we are willing to say in the book's preface, we will have learned how our *beliefs* are constrained.⁶

It is, of course, of no small importance that this slide, between talk of what we are willing to say and talk of what we believe, occurs *in the context of inquiry*. What someone is willing to say while engaged in a comedic performance—what she is willing to say when her overarching aim is to amuse, or to shock, or to provoke her audience—does not incline us in the least to ascribe beliefs to her. In contrast, when the aims to amuse, shock and provoke play a secondary role, as they do in the context of inquiry, the context in which we write and read books and papers—when her overarching aim is (to speak very roughly) to tell the truth as it pertains to the propositions under consideration—then we are inclined to regard her decisions as to what she is willing to say as revealing what it is she believes.

And what is the aim to tell the truth as it pertains to a proposition P? It can't be just the aim to tell *a* truth. If all we wanted were *a* truth, we would never risk error—we would never be willing to say anything in the context of inquiry we weren't absolutely certain is true. If all we wanted were *a* truth, "P or not-P" would satisfy us. That we *are* willing to risk error reveals that, in addition to truth, we also want content—we want to have a line on things. So, when we decide whether to say that P in the context of inquiry, we are weighing at least two competing goods: the good of saying something true, and the good of saying something contentful.

It stands very much in favor of the account just described that it seems to capture in a simple and straightforward way how we pursue the aim to tell the truth: our being unwilling to say things unless we are highly confident of their truth reflects our desire for truth, and our willingness to say things of which we are not fully confident reflects our desire for content.

The trouble is that it is *too* simple. According to this account, the propriety of a person's being willing to say that P in the context of inquiry turns entirely on the propriety of her having a confidence in P that exceeds the threshold: her being willing to say that P in the context of inquiry is beyond criticism so long (and only so long) as her confidence in P exceeds the threshold and her being that confident in P is itself beyond criticism. But this *Confidence Threshold Requirement* cannot be correct—and for two reasons.

⁶ The preface paradox was first presented in Makinson 1965. The paradox arises from the thought that it seems, on the one hand, reasonable that an author should believe, of each sentence in her ambitious book, that this sentence is true, yet write in the preface (precisely because it is reasonable for her to believe) that, given the ambition of the work, not all the sentences are true. On the other hand, if she does this, she runs flagrantly afoul of the apparently compelling requirement that a reasonable set of beliefs be consistent. Much more on all this below.

III. Why the Threshold Requirement Can't Be Correct

The first reason the Confidence Threshold Requirement cannot be correct is that it simply doesn't capture what our intellectual lives are really like. By the lights of this requirement, we might do a perfectly adequate job of determining what we should be willing to say in the context of inquiry—adequate to the point that it places the position we come to adopt beyond criticism—in a completely piecemeal way: by simply examining each proposition one at a time and determining whether we are warranted in investing in this proposition a confidence that exceeds the threshold. To be sure, there would be a probabilistic constraint and an epistemic constraint that would require us to assess as well what warrant the rest of our investments of confidence have. After all, our investment of confidence in one proposition will open our position to criticism if it does not cohere probabilistically with the rest of our investments of confidence, or if it reflects an assessment of our evidential situation that isn't fair, given how we assess our evidential situation in other cases. But there would be no call—no place—for us to evaluate propositions *qua* things we might be willing to say:⁷ no call, no way, to consider how they hang together *as things to say*.

This is just plain false to the way we actually decide what we should be willing to say in the context of inquiry. There we are looking, not for some individually probable things to say, but for a story to tell about how the world is.⁸ And we evaluate propositions by looking at how their adoption might fit in, or fail to fit in (and if so require alteration of), the story we have so far. Granted, we want our story to be true—and so we pay attention to the extent to which it, and the things that go into it, are worthy of our confidence. But (as already noted) truth is not all we want. If it were, we would be happy with a story of the world that was nothing more than a tautology. We also want our story to be contentful. It is as we assess ways we might alter our story of the world that we weigh the competing goods of saying something true, and saying something contentful.

Of course, the account under discussion—according to which we are warranted in being willing to say that P if and only if we are warranted investing confidence above the threshold in P (where that means being at least more confident than not that P)—cannot hope to capture how we decide what story of the world we should be willing to tell in the context of inquiry. Any such story will inevitably be one in whose truth we will have no business being confident. A big, ambitious, story brings with it a very great risk of error. Yet we can make sense of our

⁷ Or, more properly, to say *to be true*.

⁸ This way of putting it is due to Jonathan Roorda. See Roorda 1997, pp. 148–9.

willingness to tell such a story in the context of inquiry. What warrants our being willing to tell such a story is that it represents (and here I am supposing it is the right sort of story, not some crazy conspiracy theory) a measured and reasonable trade of the likelihood of obtaining one desideratum for the certainty of getting more of the other: a trade of the likelihood of truth for content. And such a trade is necessary, if we are going to go after such a story. For the content such a story affords us is not available for any lower price.⁹

The second reason the Confidence Threshold Requirement cannot be correct is that, precisely *because* it is incompatible with the Requirement of Deductive Cogency, the Confidence Threshold Requirement cannot account for the force *reductio* arguments exert upon us as we decide what we should be willing to say in the context of inquiry.

This is something Christensen contests. He maintains (79–96) that the force of *reductio* arguments can be fully explained without assuming that deductive cogency has any purchase on what a person ought to be willing to say in the context of inquiry. He notes quite correctly that, by the lights of the Confidence Threshold Requirement, a *reductio* of a small number of propositions can constitute a very serious criticism, showing that there is no way, consistent with a probabilistic constraint on our investments of confidence in propositions, to invest enough confidence in each of the propositions involved so that they will all exceed the confidence threshold. So, for example, the *reductio* I used to motivate the Requirement of Consistency shows that there is no way compatible with the probabilistic constraint on warranted confidence for me to have a confidence of greater than 0.67 in any one of the three propositions that serve as the *reductio*'s premises without having a confidence of less than 0.67 in at least one other of them. The only *reductios* that lack critical bite by the lights of the Confidence Threshold Requirement are those that employ very large numbers of propositions to derive a contradiction. But, maintains Christensen, such *reductios* lack force anyway.¹⁰

But this won't do. While the Confidence Threshold Requirement does thus issue the right verdict about the outrageous behavior to which I appealed by way of motivating the Requirement of

⁹ I actually think that one can make sense of our making such trade-offs (and so locate properly the role our warranted investments of confidence play in determining what it is rational for us to be willing to say in the context of inquiry) without having to assume (as I think one does not want to assume) that content is something that is subject to measure. See Kaplan 1996, pp. 121–40.

¹⁰ This is a line of reasoning that has been championed in Hawthorne and Bovens 1999, and in Weintraub 2001. A more nuanced version of this line of reasoning is to be found in Sturgeon 2008.

Consistency, it issues the wrong verdict about the outrageous behavior to which I appealed by way of motivating the Requirement of Deductive Cogency. That is because, by the lights of the Confidence Threshold Requirement, there *is* nothing outrageous about that behavior. It is entirely compatible with the probabilistic constraint on warranted confidence that, for any non-maximal threshold, my confidence in each of the three things you point to exceeds the threshold, but my confidence in their consequence, P, does not. I am thus free, by the lights of the Confidence Threshold Requirement, to dismiss a demonstration, that three things I've said (even that two things I've said) jointly entail something I am unwilling to say, as being *no criticism whatsoever* of my being willing to say them.

Further, the way in which the Confidence Threshold Requirement explains the outrageousness of the behavior it *does* recognize as outrageous will not hold up. The explanation requires us to acquiesce in the thought that, so long as a *reductio* has sufficiently many premises that we could warrantably invest in each of them a confidence that exceeds the threshold, we are free to dismiss as offering no criticism at all a *reductio* meant to criticize our being willing to say, in the context of inquiry, the things that serve as the argument's premises. But how can that be right?

Suppose that you charge me with producing a chronology of the events that preceded a serious accident.¹¹ Suppose that the chronology I produce takes the following form: "At 8:00 am event A occurred. After event A occurred, event B occurred [...]. After event Y occurred, event Z occurred. After event Z occurred, event A occurred." You ask me whether I concede that the x-occurred-after-y relation is transitive and not reflexive. Naturally, I do. "Then how," you ask, "can you possibly offer this as a chronology of the events leading up to the accident? Given what you have (rightly) conceded, your chronology logically implies a contradiction."

There can be no question but that the *reductio* you have produced has critical bite: it exposes the fundamental inadequacy of the chronology I have produced. This inadequacy is in no way mitigated by the fact that the *reductio* requires twenty-nine propositions to derive its contradiction—by the fact that the derivability of this contradiction is compatible with my being entirely warranted in having a degree of confidence greater than 0.96—greater than any reasonable threshold—in the truth of each of the propositions. The idea that only *reductios* of a relatively small number of propositions have critical bite, when it comes to what can rationally be said in the context of inquiry, simply won't wash.

¹¹ I offer the following argument in Kaplan 2002, pp. 459–60, fn 20.

Christensen's response (89–90) is that the chronology we have just imagined my having produced is “fundamentally inadequate,” but that the case is not described in enough detail to permit us to determine wherein the inadequacy lies. We would have to know a lot more about why I produced the chronology I did—and, in particular, how confident I was in each of its elements and why—before we can say what is wrong with the chronology.

But that is surely wrong. My production of the chronology opens my position to criticism on grounds of its inconsistency no matter *why* I produced it, and no matter *why* I am as confident as I am in each of its elements. It opens my position to this criticism even if, as is perfectly imaginable, I am entirely warranted in being as confident of each proposition as I am: even if the evidence is such as to warrant my finding each of the propositions enormously credible even as I recognize that they cannot all be true. That is because we expect a chronology of the events leading up to an accident to provide (at the very least) a consistent story as to what happened and when. And we expect much the same of any explanation, historical narrative, prediction, even of a simple report as to what has been observed: we expect it to provide us with, at the very least, a consistent story. The one I offered does not.

Not that it is any part of Christensen's brief that, provided I am warranted in being confident enough in each of the elements of my chronology, my willingness to say what I do is beyond criticism. While sympathetic to the Confidence Threshold Requirement, Christensen is not an advocate of the Requirement. On the contrary, he seems to think that the Requirement gets our reaction to the lottery wrong. He writes (27) that “we are reluctant to assert unqualifiedly ‘This ticket will not win’ even when the lottery is large.”¹² It is thus available to him to hold—and I read him as holding—a different view: namely, that the phenomenon of the damaging *reductio* is best understood in a piecemeal way, case by case. In this manner, the thought would be, both the apparent excesses of the Requirement of Deductive Cogency, and the apparent excesses of the Confidence Threshold Requirement can be avoided: we don't have to endorse the wholesale condemnation of inconsistency, yet we can acknowledge the bite of some *reductios* (such as your *reductio* of my chronology) that apparently employ lots of premises. In particular, we can explain why a *reductio* has bite in a particular case by appeal to features peculiar to that case—features not always present when a *reductio* of a position is available—thereby

¹² Thus the only part of the story I sketched in section II that I think Christensen would wholeheartedly endorse is that it is a condition on your being warranted in willing, in the context of inquiry, to say that P, that you be warranted in being (at the very least) more confident than not that P.

rendering the explanation of the *reductio*'s bite quite independent of the Requirement of Deductive Cogency.¹³

But this view does no better than the Confidence Threshold Requirement at making sense of what makes for a legitimate criticism of a position. On this view, for an argument to the effect that a position is inconsistent to be effective—for it to be *any criticism at all* of the position—it has to come with the right kind of commentary as to what there is about the position in question (what's wrong with the position in question) that warrants regarding the *reductio* as having critical bite. That is to say, it would (on this view) be a perfectly legitimate defense in the face of the charge that one's position is inconsistent—a perfectly legitimate ground on which to deny that the charge is *any criticism at all* of one's position—that the person who has made the charge, and established that the charge is true, has failed to establish that there is any such commentary.

But it would be outrageous to suppose that, absent such a commentary, one can shrug off a *reductio* of one's position as having no more critical bite than, say, the observation that "Monday" begins with an "M". To be sure, some *reductios* are more important than others, some more damaging than others. But then the defense against an unimportant *reductio* is to say *why* it is unimportant, or how it is easily handled, or (even if it is not easily handled) why it is worth our while to set it aside for now. It is not to act as if the *reductio* is no criticism at all.¹⁴

Why would Christensen even suggest otherwise? I think it is in part because of the way he thinks about the lottery paradox. He regards as innocent the inconsistency involved in being willing to say, of each ticket in a fair 1000-ticket lottery that (one is willing to say) will

¹³ See 93. I thank Scott Sturgeon and David Christensen for getting me to see that the latter can, and should, be read this way.

¹⁴ None of this is meant to deny that there are cases in which there is more wrong with being willing to say a particular inconsistent set of things in the context of inquiry than mere fact that the set is inconsistent. There are cases (the lottery is one) in which, as Christensen notes, there is more than one thing wrong: one also simply shouldn't be willing flatly to say, of a ticket in fair lottery, that it will lose. So I have no stake in denying Christensen's claim (91–3) to have found a pair of cases—one in which a person is willing to say an inconsistent set of things, the other just like it except that there is no inconsistency—in which the move from the first case to the second does not remove our sense that the person shouldn't be willing to say the things she is willing to say. My only stake is in denying that there is nothing in itself wrong with being willing to say an inconsistent set of things in the context of inquiry. That is to say, all I am denying is that, only if there is something *else* wrong with being willing to say those things, does a *reductio* using those things as premises constitute a criticism of her willingness to say them in the context of inquiry. (I thank an anonymous referee for pressing me to say more on this.)

produce one winner, that this ticket will lose.¹⁵ He thinks that one of the excesses of the Requirement of Deductive Cogency is that it demands that we regard the inconsistency otherwise.

But where is the innocence? To be willing to say of each ticket in a 1000-ticket lottery that it will lose, while being willing to say that the lottery is fair and so one ticket will win, is to be willing to tell a wholly incoherent—because manifestly inconsistent—story as to how the lottery will turn out: a story on which each ticket will lose but one ticket will win. I do not see how this story of how the lottery will turn out can be thought to be in any way better than the story we imagined my having produced—a story Christensen agrees is fundamentally inadequate—as to what was the chronology of events that preceded the accident.¹⁶

But if the Confidence Threshold Requirement fails so miserably at making sense of what our intellectual lives are really like, why does it tempt us all the same? The answer, I think, is that it is because we operate on the assumption that “believe” is univocal: that, whenever we think and talk about what we do and do not believe, and about what

¹⁵ Though, as we have already noted, he thinks that being willing to say, of even one ticket, that it will lose is *not* innocent; on the contrary, it is something we are reluctant to say. But he thinks there are (what he calls) lottery-like cases, which exhibit the innocent inconsistency present in the lottery case, but where the inconsistent set is made up of a set of things each of which a person *should* be happy flatly to say. And he thinks these pose a “strong” “intuitive challenge” to those who would champion deductive cogency. For example, he tells us that he believes, of each day in the coming year, that the paper will be on his doorstep by 6:30 in the morning, but thinks it irrational to believe (as he must, if his beliefs are to be deductively cogent) that there will not be a single late paper all year. See 54–5.

¹⁶ It might be thought that the maneuver just executed is not always available, because not every set of things a person is willing to say in the context of inquiry tells a story. (Indeed it might be thought that the Requirement of Deductive Cogency is to be rejected precisely because the set of things you are willing to say in the context of inquiry really doesn’t constitute a story.) It may be no accident that Christensen prefers, when examining what claim deductive cogency has on a rational inquirer, to focus on sets of things an inquirer might write that don’t seem to add up to stories—for example, a telephone directory (35n) or a report on the outcomes of a series of horse races (92–3). But notice that, even in these cases, we have stories being told—as to how many people whose names are listed in the telephone directory actually have numbers ending with an even digit; what percentage of the numbers listed in the directory actually belong to people with Hispanic surnames; what percentage of the horses from a given stable failed to finish in the money; and so on. And, even in these cases, an inquirer’s willingness to say the things she has written is open to criticism in the context of inquiry if those things jointly entail something that she is unwilling to say. So, for example, her reporting of the outcome of each one of a substantial number of horse races opens her position to criticism in the context of inquiry if there is a published summary report, as to what percentage of horses from each stable have finished in the money, that is not consistent with what she has reported as to the outcomes of the races, and she is not prepared (because, let us suppose, the report has been issued by an extremely reputable outfit) to say that the published summary report is mistaken.

we should and should not believe, it is one and the same thing that we are thinking and talking about. So, since we clearly think and talk of belief both as a state of high confidence and as a willingness to say, in the context of inquiry, that something is so, we conclude that that we are warranted in being willing to say something in the context of inquiry just if we are warranted in being highly confident that it is true.

I think that one moral to draw, from the failure of the Confidence Threshold Requirement to capture what our intellectual lives are like, is that it is a mistake to think “believe” is univocal. We quite legitimately use “believe” to talk of what we are highly confident is so, and to talk of what we are willing to say in the context of inquiry. But we make a mistake if we do not recognize that, in so doing, we are using “believe” to talk about two distinct things.¹⁷

Indeed it is this very mistake that, I submit, is responsible for the long career of the lottery paradox. We are clearly warranted in being extremely confident, with respect to each ticket in the lottery, that it will lose; we have taken that as meaning that we are warranted in believing of each ticket that it will lose; we have taken that, in turn, to require us to be willing, in the context of inquiry, to say of each ticket that it will lose; and we have rightly been disturbed by the fact that this last (together with our being willing to say that one will win) commits us to being willing to say an inconsistent set of things—to a position that is *eo ipso* open to criticism (as the Requirement of Deductive Cogency rightly says).

But, once we recognize the way in which our use of “believe” fails to be univocal, there is no cause here to be disturbed. True, high confidence is called upon to obey a probabilistic constraint that permits us, without fear of thereby opening our position to criticism, to invest high confidence in each member of an inconsistent set of propositions. True, the willingness to say something is so in the context of inquiry is called upon to obey the Requirement of Deductive Cogency, which says that we open our position to criticism if we are willing to say true in the context of inquiry each member of an inconsistent set of propositions. But once we recognize that *what we are highly confident is the case* is one thing, *what we are willing in the context of inquiry to say is the case* is another, and that our being warranted in being highly confident that

¹⁷ This represents an alternative to my position in Kaplan 1996, where (having tacitly assumed that our use of “believe” is univocal) I argued that, since a person’s set of beliefs is open to criticism unless deductively cogent, but a person may, without thereby opening her position to criticism, have any high degree of confidence less than 1 in a set of propositions that is *not* deductively cogent, belief simply cannot be construed as a state of confidence (see chapter 3); and where I suggested (p. 154) that, because, on our ordinary conception of belief, belief is just a state of confidence above a threshold, we cannot but conclude that our ordinary conception of belief is incoherent.

P does not determine of itself whether we are, in the context of inquiry, warranted in being willing to say that P, we no longer have here the makings of any conflict.¹⁸ Paradox dissolved.

IV. The Preface: Saying It's All True

Having said why I think we cannot make sense of our intellectual lives without assuming the propriety of the Requirement of Deductive Cogency, I want now to take up the challenge of saying how we can possibly make sense of our intellectual lives if we *do* assume the propriety of the Requirement of Deductive Cogency. How can it be that a person opens her position to criticism unless, upon completing a long, ambitious book of historical scholarship, she is prepared to say that everything she has said in the book is true?

I want to begin my answer to this question by posing another. If our author shouldn't be willing to say that everything she's said in her book is true, what *should* she want to say about the matter as to whether her book has errors? Should she be willing to say, "This book contains errors"? Christensen admits (35) that, despite the fact that (as she must concede) it is overwhelmingly probable that she has made some mistake in the body of her book, we would find it strange if our author flatly said in the preface, "This book contains errors." He suggests (43–4) that this may be because (as some writers have claimed) we understand unqualified assertions to express knowledge, not merely belief.¹⁹ And so, he suggests, the awkwardness we see in her saying flatly that her book contains errors may be explained by our thinking that this is not something she counts as knowing.

But this explanation won't do. We don't regard it as a condition on the appropriateness of her saying each of the *rest* of the things she flatly says in her book that it count as something she knows. And if we don't in general demand that an author of a scholarly book say nothing in her book that she cannot reasonably count herself as knowing to be true, why would we impose this condition in this one instance?

I want to suggest another explanation. If our author flatly said in her preface, "This book contains errors," she would invite the reaction: "Then why didn't you correct the errors before the book went to press?" And why would she invite that reaction? Because we all know how

¹⁸ This raises, of course, the question as to what exactly *is* the role that our warranted investment of confidence in propositions plays in determining which of those propositions we should be willing to say true in the context of inquiry. I offer a sketch of an answer in Kaplan 1996, pp. 121–40.

¹⁹ Christensen cites Williamson 2000, ch. 11. There is by now a rather considerable literature on whether (and, if so, in what sense) knowledge is the norm of assertion.

improbable it is that any book of any reasonable ambition will be error-free. But we also know that no reasonably ambitious book—indeed no contentful story of how the world is—is available at any lower price. It is a fact of life that one cannot get a contentful story of how things go in some part of the world without, *eo ipso*, getting a story that is likely false. We expect an author to have recognized this fact of life, to have come to terms with it, and (having written her book) to stand behind the book she has written. Why tell the story she has told in her book, if she is not going to stand behind it? Why tell the story if it is not going to be *her* story? But to stand behind the story she has told in her book just *is* to be willing to say that the story—that everything she has said in the book—is true.

This last claim may seem an exaggeration. But is it really? Suppose, on being asked whether everything she's said in her book is true, she doesn't say, "Yes." Instead, she declares herself to be on the fence: unwilling to say that everything she's said in the book is true, and also unwilling to deny it either. Is there any question but that we will read her as thereby distancing herself from the story she's told in her book, as denying that it really *is* (or is anymore) part of her story as to how the world is? If she isn't willing to stand by everything the story says—if she isn't willing to say that everything in the story is true—it is not her story.

Christensen would, I expect, deny this last claim. His view appears to be (77) that it is entirely compatible, with a person's being suitably committed to the sort ambitious story our author has told in her book, that she be unwilling to say that her story is true. She need only be willing to say that it is *approximately true*. He sees nothing to be gained by making the stronger, immodest, claim of truth—he doesn't see what difference it makes whether one is willing to make the stronger claim. Perhaps that is because, on Christensen's view, being suitably committed to a story is entirely compatible with taking only limited responsibility for what that story's consequences are—with being indifferent to whether the story is even consistent. "Suppose, for example," Christensen writes (93), "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?" Christensen argues that it would not. After all, he writes (94), the author is already warranted in "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."²⁰

²⁰ The emphasis is Christensen's.

But, while her having the latter sort of commitment to the story she's told in her book is indeed compatible with our author's being willing to say no more than that the story is approximately true, it falls far short of what is required for her to count as standing behind her story—for the story to count as being *hers*. After all, it is entirely compatible with her being willing to say that the story she's told in her book is approximately true, that (far from standing behind the story) she has already rejected the story in favor of another that (by her current lights) preserves the many things the story in her book gets right, and gets right some of the things the story in her book gets wrong.²¹

None of this is to say, of course, that our author should stick doggedly by her story in its every detail no matter what reason she may come to have for thinking that the book contains a mistake. On the contrary, presented with new evidence that contradicts one of the claims she has made in her book, she should admit her error. That is, she should stand by her book—she should not shrink from saying that everything in it is true²²—until she has found special reason to think it contains an error. And what makes something a special reason to think the book contains an error? It seems to me that we can safely say this much. We find it perfectly appropriate that an author should stand behind her book even in the face of the high probability that it, on account of its being a book of significant ambition, contains some error—even if there are other authors of comparably ambitious works, no less careful (and, perhaps, some more meticulous) than she, who have nonetheless failed to write error-free books. Those aren't special reasons for not standing behind her book; they are reasons that are routinely present—some of them reasons that are *always* present—when an author (perfectly appropriately) stands behind an ambitious book. A special reason for not standing behind her book—a special reason for thinking it contains an error—must (at the very least) be a reason an author of an ambitious book it would be perfectly appropriate to stand behind *doesn't* routinely have for thinking her book contains a mistake.

Thus it is that, if our author flatly said in her preface, “This book contains errors,” our reaction would be: “Then why didn't you correct the errors before the book went to press?” Our expectation is that an author will flatly say that there is an error in her book—she will stop standing behind the book she has written—only when she has a special

²¹ Indeed it is often said to be the mark of a progressive research program that theories produced earlier in the history of the program, and that have since been rejected, can nonetheless be said, by the lights of their successor theories, to be approximately true.

²² Though, of course, she has no particular obligation say it is true in her preface either.

reason to suppose the book contains one. And, our thought is, if she has such a special reason to think something she's said is in error, she should *either* have gone to the parts of her book that this reason gives her ground to think responsible for her error(s), and done what is necessary to remove them, *or* explain why she hasn't.²³ As things stand, she has no special reason—no reason one doesn't routinely (if not *always*) have with respect to a book of any significant ambition that an author might perfectly appropriately stand behind—to think her book contains errors. And that is why she shouldn't be willing to say—and we would find it odd if she did say—that it does.

V. The Preface: A Matter of Modesty

But shouldn't we worry that, as Christensen maintains, our author would be guilty of an unreasonable immodesty were she willing to say that the story she's told in her book is true? I don't see why. After all, our author is still free to engage in the graceful practice of exculpating those whose help she has acknowledged from such errors as remain. She is even free to confess to being confident that her book will turn out to contain some error or other. The only thing she cannot do is flatly take back what she's said in the rest of the book: she cannot flatly say (or say anything that entails) that her book has errors.²⁴ But this prohibition does not in any way stand in the way of our author's displaying her modesty. As we saw, we wouldn't *interpret* her flatly saying her book contains errors as an expression of modesty. We would interpret it as a confession of intellectual sloth.

Might this not require her (as Christensen claims it does, 46–8, 104) to be immodest in another way—to take the position that her book is unique in being error-free? Not at all. For that to be so, the

²³ Some possible explanations: it would be hideously expensive to find and correct the errors; the book is just a preliminary report and these errors will be corrected with all the others that are pointed out before the final report is written; the book is meant simply to float a story that will provoke criticism and spur research, and it does not matter terribly much if there are errors in it.

²⁴ Christensen seems to think (36) that prefatory remarks often violate this prohibition. He cites a prefatory remark of David Mellor's (36n), in which Mellor thanks those who have offered him stimulation and then adds, "The mistakes are all mine." I am inclined to read this as no admission of error, but rather as "The mistakes, if there are any, are mine." A desperate reading? I think not. Suppose you said, "If there is a flood, the responsibilities are divided as follows: the Fire Department takes responsibility for the injured, the Coroner's Office takes responsibility for the dead." And suppose I objected, "But not every flood involves loss of life!" Wouldn't you be perfectly within your rights to respond that you never said otherwise? Wouldn't you be perfectly within your rights to expect me to have understood what you said about the responsibilities of the Coroner's Office in this way: "The Coroner's Office takes responsibility for the dead, *if there are any*"?

Requirement of Deductive Cogency would have to commit her to being willing to say, of every book other than her own, that it contains mistakes. But the Requirement does no such thing. On the contrary, the very consideration that speaks against her flatly saying that her own book contains errors (she has no special reason, no reason one doesn't *routinely*, if not *always*, have with respect to a book of any significant ambition that it would be perfectly appropriate for an author to stand behind, to think her book contains errors) speaks just as strongly against saying, of some book she *hasn't* written (and, perhaps, hasn't even read), that it contains errors. The inappropriateness of our author's flatly saying that her book contains errors doesn't just come from its being *her* book; it comes also from the inappropriateness of flatly saying, of *any* book, "This book contains errors," if all one has to support that claim is the fact that the book is ambitious and so it is highly probable that it is mistaken in some detail. If you say flatly, of a particular book, *any* book, that this book contains errors, we expect you to have some better reason than this—some special reason—for saying so.²⁵

VI. The Preface: Saying the Improbable

But how is our author supposed to square her willingness to say her book is error-free with the fact that it is overwhelmingly probable that the book is *not* error-free? How can our author be willing to say, in the context of inquiry, "My book contains no errors but, mind you, I regard it as overwhelming probable that I am mistaken in this"? Isn't that a completely bizarre thing to say? I imagine Christensen would think so. He writes (45) that, were an author to say that he *believes* his book contains no errors, but that it is overwhelmingly probable that he's mistaken in this, his performance would leave most people incredulous, "if," he adds, "they took [it] as something other than a deliberate joke at his own expense."

Let's take a closer look. We have already seen that the following principle won't do: you should be willing to say that P in the context of inquiry if you are aware that P is highly probable. Take any fat

²⁵ Of course, there is an important respect in which our author will not treat her book as equal to a book written by someone else—or, at least, to a book written by someone else, that she hasn't read and found to contain (as she wrote her own so it would contain) nothing she isn't willing to say in the context of inquiry. She will be willing to say of her own book, and of any book that she has read and agreed with in the sense just described, that it contains no errors. That is something she won't be willing to say—but, of course, it is something we would not *expect* her to be willing to say—of a book she hasn't read, or has read but in which she has found written things she is not herself willing to say in the context of inquiry.

book of history off a library shelf. You are aware that it is highly probable that not everything its author says in the book is true. Yet, as we just saw, it would be appalling for you flatly to say, never having even opened the book, “Not everything in the book is true.” Only if you have a special reason to think the book contains some error—a reason you don’t have for *every* book of any significant ambition—are you entitled flatly to say that the book contains errors.

Now consider the following principle: you should *not* be willing to say that P in the context of inquiry if you are aware that P is *improbable*. It is this principle that Christensen is implicitly invoking when he argues that our author should not be willing to say that everything she’s said in her book is true. It is the violation of this principle that he holds responsible for the apparent bizarreness of our author’s saying, “My book contains no errors but, mind you, I regard it as overwhelming probable that I am mistaken in this.” Is *this* principle correct?

There is no question but that there is an important class of cases in which the principle holds true. These are the cases in which the improbability you attribute to P, your confidence that P is false, derives from something more than the mere ambition of P. It derives from the conviction that there is some defect in the argument that has been made (or is available to be made) for P—a defect that isn’t routinely (if not *always*) present in the arguments available to be made for propositions of comparable ambition, a defect that one could (perfectly legitimately and without fear of thereby having done something quite outrageous) cite in the context of inquiry by way of issuing a critique of our author’s having said that P. In this important class of cases, the cases in which (we can say) there is a special reason to suppose that P is false, the principle—that you shouldn’t be willing to say that P in the context of inquiry if you are aware that P is improbable—is correct.

But what about the remaining cases? What about the cases in which there is no special reason to suppose that P is false, there is no defect in the argument that has been made for P? What about the cases in which the improbability you attribute to P derives simply from the ambition of P? Does the principle hold true in these cases?

It seems to me that it does not. You can take any fat book of history off the shelf and, without even opening it, say (with good reason) that it is extremely improbable that the story that book tells is true. Yet we regard this as no criticism whatsoever of the book, no mark against the story the book tells, no sign of there being any impropriety in the author’s being willing to tell that story, no impropriety in its *being* her story. But, as I have already argued, for it to *be* her story, she has to be willing to say the story is true. And that is just to say that, when the improbability of P derives solely from its ambition, the fact that P is improbable—the fact

that you are aware that P is improbable—offers no reason for you not to be willing to say that P. In these cases, the improbability of P is simply not germane to the propriety of being willing to say that P.

My suggestion is that the apparent bizarreness of our author's being willing say, "My book contains no errors but, mind you, I regard it as overwhelmingly probable that I am mistaken in this," arises from a failure to distinguish the two sorts of cases. It arises from supposing that, whenever she says, of some proposition P, "I regard it as overwhelming probable that not-P," she must be saying something germane to the propriety of her being willing to say that P—she must be confessing to having a special reason to think that P is false. The supposition is false.

Indeed, we can see that the supposition is false by seeing that our author can, by adding a suitable commentary as to *why* she is saying what she is, erase the specter of bizarreness from her willingness to say it. Our author can say this: "I stand by my book: ask me if everything I've said in the book is true, and my answer is, 'Yes'. Am I *confident* that everything I've said in the book is true? Of course not. It is in the very nature of a book of significant ambition that it most probably contains some mistake or other. But I don't *have* to be confident that everything I've said in my book is true to be entitled to stand behind my book—to be entitled to say that everything I said in it is true. All that is required is that the story I have told in my book represents a reasonable and measured trade of likelihood of truth for content. And, having written the book with due care, I am satisfied it meets this constraint. Of course, I am prepared to alter my stance should I come upon a special reason to suppose that I have made a mistake. But the fact that the book is ambitious and so likely to contain some mistake is not such a reason."

But doesn't the fact that such an elaborate commentary is required to explain away the apparently difficult consequence our author must face should she commit herself to the Requirement of Deductive Cogency—our author must be prepared to say, "There is no error in my book, but, mind you, I am confident that I am mistaken in this"—just go to show how grievous that consequence is? I don't think so. After all, it isn't as if the unadorned, "My book contains errors," is something our author can comfortably say. Remember, were our author flatly to say, "My book contains errors," we would not think that she was confessing to an admirable modesty. We would think that she was confessing to intellectual sloth—to being aware of errors in her book that she hadn't bothered to correct.

Christensen thinks it is easy enough for our author to avoid fostering this impression, yet still say something that entails that her book contains errors. She need only say something like, "In time, errors will be found in this book." He writes (36) that, if asked the question,

“Will errors be found in your book?” there is nothing at all unnatural about her saying, simply, ‘Yes.’” But is this really so?

Imagine that we continue the conversation with our author. “Don’t you really mean to be saying,” we ask, “not that errors *will* be found in your book, but that errors *will most probably* be found in your book?” She demurs. “No, I am saying that errors *will* be found in my book.” We press on. “So you are not saying just that there are very probably errors in your book?” “No, I’m not,” she replies. “You are saying, flat out, that there are actual errors in your book?” “Yes, that’s precisely what I’m saying,” she replies. Is there any question but that, once the conversation has gone this far, we will be thinking (or, at the very least, suspecting very strongly) that our author is, in effect, telling us that she has a special reason (a reason that doesn’t derive from the mere ambition of her work) for thinking that her book contains errors? Won’t we, once again, be thinking that she is not confessing to modesty; she is confessing to intellectual sloth?²⁶

I do not mean to suggest that there is *nothing* our author could say by way of providing a commentary that will defeat the expectation that would be fostered by her flatly saying, “This book contains errors”—the expectation that she has a special reason for thinking her book contains at least some error. On the contrary, it seems to me she could say something like this: “I am not aware of any errors in this book. But it is all but inevitable that a book of this ambition will have errors. And so I have no choice but to say that this book, like any other book of similar ambition, has errors.”²⁷

Rather, what I am suggesting is this. There is nothing uncomplicated our author can flatly say as to the matter of whether her book contains errors. Whatever she says will, on pain of conveying some message she does not mean it to convey, require some accompanying commentary that explains why she is saying what she is. And there is a perfectly good commentary available to her that will explain how she could be prepared to say her book is error free, even though she is very confident that it is not. The availability of that commentary, I suggest, gives the lie to the commentary reproduced at the end of the last paragraph—the commentary that would accompany her saying that her book contains errors. It gives the lie to the thought that she has no choice but to say that her book contains errors.

²⁶ I should note that, in imagining this dialogue, I am following Christensen, who uses an imagined dialogue with an author who *denies* that his book contains errors (45) to bring out what Christensen takes to be the absurdity of an author’s being willing to say that her book is error-free.

²⁷ I thank Scott Sturgeon for suggesting as much.

VII. A Worry about Downstream Consequences

But Christensen thinks that our author is going to have a lot more explaining to do, if she decides she is prepared to conform to the Requirement of Deductive Cogency and say that her book is error-free: her decision is going to have significant downstream consequences—both for what she is going to have to be willing to say is the case and for what she is going to have to be willing to say she should do—that (even if we credit a use of “belief” to talk of what one is willing to say in the context of inquiry) are manifestly counterintuitive.

Let’s begin with the first sort of consequence. Christensen imagines (49–52) that our author is rational to believe that, if her book is error-free, then either she is extremely meticulous or she is amazingly lucky. She knows that she is not that meticulous. So, given that she believes her book is error-free, deductive cogency demands that she believe that she is amazingly lucky. He also imagines that she is rational to believe that, if anyone wrote an error-free book in her field, “it would soon lead to adulatory reviews, prestigious speaking engagements, and opportunities for professional advancement.” So, given that she believes that her book *is* error-free, deductive cogency requires her to believe that she will soon receive adulatory reviews, prestigious speaking engagements, and opportunities for professional advancement. For similar reasons, deductive cogency requires our author to believe that her colleagues in the field, Professors Y and Z, who (it is mutually acknowledged) are far more meticulous than she, are in for a big surprise when it becomes apparent that her book is error-free (as none of theirs have proved to be); that she will win the large monetary prize given (by an organization dedicated to finding errors in books in her field) to authors of books in the field that last a year in publication without any error’s being spotted; that (in the rational belief that, if she wins the prize, she will buy an Alfa-Romeo) she will a year from now be behind the wheel of an Alfa.

Christensen thinks that all this needs to be taken into account before we tally how well a view that endorses deductive cogency handles the preface. And he thinks that, once it *is* taken into account, the tally will not be favorable to the advocates of deductive cogency. “[T]he intuitive irrationality of these beliefs,” he writes (101), “[...] makes them *prima facie* counterexamples to cogency requirements.”

He appears to think no differently about the second sort of downstream consequences he thinks accounts that endorse deductive cogency have—consequences for what our author will have to believe about what she should do. He imagines (102) that, reasonably enough in his view, she believes (*i*) “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.”

Believing she has at present a perfectly decent car, and that when, in one year, she wins the prize for her book, she will immediately spend the proceeds on an Alfa, deductive cogency requires her to believe (ii) “I should not buy the new car now.” This despite the fact that she has just been offered an outstanding deal on a new car, and that were she not forced by deductive cogency to believe that her book is error-free, she would never believe the improbable proposition that she will be buying an Alfa in a year.

Worse still, Christensen thinks, deductive cogency can force our author to contradictory beliefs about whether she should take the offer. Our author believes, and entirely reasonably, (iii) “It’s very *unlikely* that I’ll be able to afford an Alfa in the next few years.” And if she believes, as Christensen thinks (102) she should, (iv) “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,” she must, if she is to obey deductive cogency, believe (v) “I should buy the new car now.” This despite the fact that she believes (ii), which is the negation of (v).

VIII. The Worry about Downstream Consequences Addressed

How troubling is all this for the Requirement of Deductive Cogency? The answer is straightforward when it comes to the alleged consequences the Requirement has for what our author must be willing to say she should do: not troubling at all. The only reason Christensen thinks that the Requirement forces our author to believe (ii) is because he thinks she should believe (i). But (i) is clearly false: whether it is true that our author has a decent car and will be buying a new car in a year (that is to say, whether things will, in fact, work out better for her if she doesn’t buy a new car now) has no implication whatsoever for whether she should buy a new car now.

Suppose she has taken her car to a reputable garage and they have told her that it is a death trap liable to burst into flames and explode at any moment. And suppose she has accordingly decided to buy the new car now. It turns out, later, that the garage confused her car with another: her car was actually perfectly fine. Who wants to say that she has done something she shouldn’t have? Or (to take the circumstance in which the author actually finds herself) suppose our author, though committed to buying an Alfa if (and only if) she wins the prize, is warranted in being extremely confident that she will not win, and so takes advantage of the offer on hand and buys a new car now. But suppose a year on the improbable happens: she wins the prize. Who wants to say that, in having bought a new car when she did (in having behaved as if her winning the prize was just as improbable as it was in fact), she has done something she shouldn’t have done?

Here it seems to me that the advocates of Bayesian decision analysis²⁸ have it right: what determines what our author should do is (a) how confident she is warranted in being that she has a perfectly decent car and will be buying a new car in a year, and (b) how valuable are the consequences that she regards as possible outcomes of the options amongst which she has to choose. And they are right that it is compatible with her having done what she should that, as a matter of fact, what she has done will work out worse than one or more of her alternatives would have if chosen; and that it is compatible with her having done what she should that she is warranted in being extremely confident that what she has done will work out worse than one or more of her alternatives would if chosen.²⁹

Let's turn, then, to the other downstream consequences—to the five improbabilities that, Christensen points out (quite rightly, this time),

²⁸ The classic presentations of the Bayesian decision analysis are to be found in Savage 1972 and Jeffrey 1983. See Maher 1993 for an excellent exposition and defense of a generalized (that is to say, weaker and so, at least to my mind, more defensible) version of Savage's view.

²⁹ Suppose you have the choice between (A) \$5 and (B) a gamble that will give you \$10 if P, and x if not-P. Even if you are warranted in being highly confident that P (and so highly confident that things will work out better if you choose B), you will be entirely rational to choose option A if x is dire enough—if, say, x consists in the slaughter of hundreds of people. Bayesian decision analysis gets this right, holding that you ought to prefer B to A if and only if the confidence you are warranted in investing in P exceeds the confidence you would need to invest in P for you to be (entitled to be) willing to choose B rather than A, given what x is. How much confidence you need to be warranted in investing in P before you can reasonably choose B over A is a function of how dire x is.

For reasons similar to the ones just rehearsed, we can also dismiss Christensen's complaint (104) that, once our author is willing to say in the context of inquiry that everything she's said in her book is true, she must be willing to say to a graduate student that it's a waste of time for him to look for even small errors in the book, for there are none to be found. The complaint is fair only on the assumption that it is the case—and so our author should be willing to say that it is the case—that, if there is no error to be found in the book, then the student ought not to spend any time looking for an error in the book. But, of course, nothing of the sort is the case. What determines whether it is rational for the student to spend some time looking for errors is not whether there are any errors there. What determines whether it is rational for him to look for errors is how confident he is warranted in being that there are any, how valuable he ought to regard the benefit he will realize if he finds any, and how dire the cost he will bear if he looks for some. It is, thus, perfectly compatible with her being willing to say her book is error-free, that our author should be willing to say that, in light of the professional benefit the student stands to gain by spotting an error (and publishing his results), the cost (in expenditure of time and effort, and delay of other projects) required to look for one, and the confidence he is warranted in having (and she has) that errors exist, he ought to spend some time looking for errors in the book. Indeed, even if he were proposing to search in a smaller portion of the book, that she had reason to be *confident* was error-free (and so she was confident that he will not find errors in that portion of the book), it might still (in light of the professional benefit to be gained) be reasonable for her to be willing to say that the student ought to spend some time looking for an error.

the Requirement of Deductive Cogency requires our author to be willing to say in the context of inquiry: that she's amazingly lucky, that her colleagues are in for a big surprise, that she'll win the plaudits and prize and (with the proceeds) buy the Alfa. Is it really irrational for her to be willing to say these five improbabilities in the context of inquiry?

To be sure, if said without further explanation, the five things Christensen has pointed out that our author must be willing to say—that she's amazingly lucky, that her colleagues are in for a big surprise, that she'll win the plaudits and prize and (with the proceeds) buy the Alfa—would strike us irrational to say, and would (as Christensen points out, 104) be bound to mislead in a context in which it is understood that she means to be telling the truth about the matters at hand. And, to be sure, that would be troubling if its being rational to be willing to say something in the context of inquiry required being willing to say that thing on its own, without any explanation at all of why one is willing to say it.

But, of course, no such requirement holds. After all, who among us hasn't left out of a paper or presentation something she wanted to say, was perfectly willing to say, and thought herself perfectly rational to be willing to say, *precisely because* she was only willing to say it if she had more space/time available to explain *why* she was willing to say it? It is fair to say that Christensen's five improbabilities require the same treatment: if called upon to manifest her willingness to say them in the context of inquiry, our author should not be willing to do so unless she has the time or space to explain why she is willing to say them. But, provided her explanation includes an acknowledgement that the improbabilities to which she is giving voice *are* very improbable, no one is going to be misled by our author's willingness to say them.³⁰

³⁰ But if our author were to offer such an explanation so as not to mislead any of her readers, wouldn't she be doing so at the cost of no longer flatly saying any of the things in question? Wouldn't she now count as probabilistically qualifying what she is saying? And, if so, wouldn't she thus violate the Requirement of Deductive Cogency—by refusing flatly to say that everything in her book is true, that she will win the prize, and so on? Not at all. For her to refuse to say flatly that (for example) she will win the prize, would be for her to refuse to say anything more than how very confident she is that she will not win the prize (or how probable it is that she will not win the prize). It would be to adopt a position that would in no way be opened to criticism—it would, rather, be supported—by the acquisition of a special reason (as opposed the reason she now has: her book is ambitious, and she wins the prize only if everything in the book survives the scrutiny to which it will be put by the prize-giving agency) to be precisely that confident (to regard it as precisely that probable) that she will not win the prize. In contrast, if, in addition to saying how confident she is that she will not win the prize, she is also flatly saying that that she will win it (as we are imagining she is), then she is adopting a different position. In particular, she is saying something that *is* opened to criticism by acquisition of a special reason to be very confident that her book is not error-free.

Of course, there is still the matter of how the rest of her explanation is supposed to go. But by now, that should not be so very hard to see. It is, I have argued, a condition on something's counting as our author's story as to how the world is that she be willing to say that everything in it is true. That much granted, what Christensen has shown is that it is a condition on the story she's told in her book counting as part of her story as to how the world is that she be willing to say the five improbabilities. As we saw earlier, the mere fact that its being her story requires her to be willing to say the improbable is not itself a reason for her not to be willing to say that everything in her story is true. What makes a story something worthy of being her story is its representing a reasonable and measured trade of likelihood of truth for content. The negotiation between these two goods is not properly made proposition by proposition (as the Confidence Threshold Requirement would have it), but at the corporate level. And it is compatible with this negotiation having been properly made (and, indeed, necessary if her book is to be one of any ambition) that there be propositions (such as the proposition that everything said in the book is true) that she be willing (in the context of inquiry) to say are true even though she is confident they are false.

Does this mean that she can rationally be insouciant in the face of any and all improbabilities she finds that her story entails? Of course not. Suppose she learns that, among her story's consequences, is an improbability that other stories of comparable ambition do not have as a consequence. (Perhaps this improbability is a consequence of a part of her story for which there are rival stories of comparable ambition none of which have this improbability as a consequence.) She has learned something that certainly tells against the propriety of her story—something that opens her story to criticism.

But we have to distinguish cases of this sort—cases in which the fact that a story has an improbability as a consequence constitutes a special reason to suppose that the story is false—from the case at hand, in which (we can imagine) our author discovers that her story entails Christensen's five improbabilities. The fact that her story has *these* improbabilities as consequences is no special reason to suppose that her story is false: *any* contentful story about how the world is she might have chosen to tell would have the selfsame improbabilities (or ones very much like them) as consequences. It is thus no mark against the story she has actually chosen to tell that it has them as consequences—no mark against her being willing to tell that story—and so no mark against her being willing to say true the improbabilities. Having to be willing to say these sorts of things is simply one of the costs

of being in the business of telling a contentful story about how the world is. It comes with the territory.³¹

I suspect, however, that Christensen will find little solace in that thought. On his view, thinking about the downstream consequences of a commitment to deductive cogency (104) “reveals some strong reasons for doubting the importance of cogency-respecting [categorical] belief.” “To put the point another way,” he writes three sentences later (105), “examples like those just considered above raise the following sort of question: what point would there be to a practice of selecting a favored set of propositions to ‘believe,’ if this set of propositions included propositions of the sort [our author] is required to believe by deductive cogency?”

But the question so posed strikes me as an odd one. What is under discussion is not the advisability of adopting some new practice in which we might possibly engage. It is, rather, how best to make sense of an *actual* practice in which we engage—a practice in which Christensen engaged in writing his book, a practice in which I am engaged as I write these words: the practice of flatly saying things to be so in the context of inquiry. What is the point of the practice of flatly saying things to be so in the context of inquiry? What determines the propriety of a person’s being willing flatly to say something is so in the context of inquiry? What constitutes a legitimate criticism of a person’s being willing flatly to say

³¹ But isn’t there the following concern? Each of the things our author is thus committed to being willing to say seems to suffer from an improbability that cannot be said to attach to every proposition of its type. Consider, for example, the proposition that our author will be driving an Alfa a year from now. It is improbable in a way that the proposition, that I’ll be driving a 1997 BMW a year from now, is not. (I am driving one now and I have excellent grounds for being confident that I will be doing so a year from now.) But doesn’t that mean that there is a special reason to think that she won’t be driving an Alfa a year from now? And isn’t that enough to render the improbability, that attaches to the proposition that she will be driving one, germane to assessing the propriety of her being willing to say so—and fatal to the position, which I am advancing here, that she *should* be willing to say so? Only on the assumption that, for these purposes, the proposition that our author will be driving an Alfa a year from now, and the proposition that I’ll be driving a 1997 BMW a year from now, *are* of the same type. But they’re not. While both are predictions about what car someone will be driving a year from now, the grounds upon which the predictions are being made are (obviously) quite different. And recall: a special reason to think that she will not be driving an Alfa a year from now would have to take the form of a reason to suppose something is defective in the argument she has available for the proposition that she will be driving an Alfa. It will not suffice, to provide such a reason, to note that there are propositions about what car a person will be driving a year from now for which a better argument can be provided. After all, if that *did* suffice, then no one could legitimately say in the context of inquiry something on the basis of having made a good inductive case for its being so, so long as it is something she might have (and others have) observed directly is so: the inductive case she has made is not going to be as good as the case others can make for the same thing on the basis of direct observation.

something is so in the context of inquiry? My brief has been that we cannot make sense of this practice—we cannot offer anything like satisfactory answers to the foregoing questions—unless we grant the propriety of the Requirement of Deductive Cogency.

What, then, about Christensen's misgivings about how important the practice could possibly be once we grant the propriety of the Requirement of Deductive Cogency? His doubts about the importance of cogency-respecting categorical belief, he tells us (105), don't just arise from his conviction that "such belief is determined by factors insufficiently related to truth." They arise as well from "the difficulty of connecting this sort of belief in any intuitive way with the rest of the agent's concerns, attitudes, or practices."

Christensen is surely right to this extent: the attitude of being willing to say that something is so in the context of inquiry doesn't connect with our emotional life, and with our thinking about what we should do, in the way that we (or, at least, many of us philosophers) find it intuitive to think beliefs do. It is the confidence we invest in propositions that fills this role. And, if Bayesian decision analysis is anywhere near right, there is no question but that what confidence we invest in propositions has a much more important, and much more central, role to play in our lives than what we are willing to say in the context of inquiry. The latter matters only in the context of inquiry. The former matters in every context, practical or intellectual, to which any of the propositions in question are germane. But it is one thing to say that what we are willing to say in the context of inquiry hasn't the importance of our investments of confidence, and another to suggest, as Christensen seems to be suggesting, that it has no importance at all.

Indeed, to my way of thinking, he has already voted the other way. Consider where Christensen's dismissive remarks appear: in a book—a monograph he has written—that tells a 178-page story about what claim logic has on the opinions of a rational inquirer. In the book, he takes very seriously what other people have written on the matter and the propriety of their having written what they have. Like everything else Christensen has published, his book is careful and well-crafted. It bears none of the marks (familiar enough) of work written without much worry about whether it hangs together as something to say, none of the marks (again, familiar enough) of the indifference to inconsistency to which Christensen insists he is committed.³² Christensen talks

³² Recall that Christensen holds that, in view of the fact that our author is already warranted in being highly confident that there is an error in the body of her book, it is hard to see why the slight elevation of that confidence that would result from her discovering an inconsistency in the body of the book should be alarming or even very disturbing.

the talk of the unimportance of what one is willing to say in the context of inquiry, of the unimportance of worrying about the cogency of what one is willing to say in the context of inquiry. But (and those of us who have been learning from his papers and his book have reason to be thankful for this) he does not walk the walk.

IX. Coming to Terms with our Human Fallibility

I recognize that some will be tempted to think all the same that, precisely because the Requirement of Deductive Cogency concerns itself only with the epistemic propriety of what you are willing to say in the (rather rarified) context of inquiry, the matter that has preoccupied me in this essay must count as quite peripheral to our real epistemological concerns. But I don't see it that way.

The question as to what we ought in the context of inquiry to be willing to say is so (in the sense that we open our position to criticism in that context unless we are willing to say it is so) is one that, at least for some of us, is *very* important—and very central—to our ordinary lives. The amount of time and energy we devote to the production and scrutiny of books, articles and lectures and the propriety of their authors' willingness to say what they have said in them, testifies to our preoccupation with the with assessing the epistemic propriety of being willing to say, in the context of inquiry, that something is so. It is no wonder. Questions about whether we should (in the foregoing sense) be willing to say that P in the context of inquiry are questions about what our line on things should be, what our position on things should be. Not only are these questions ones that we have always thought an inquiry into the warrant for belief should ultimately answer, they are questions that lie at the heart of epistemology—whose goal, after all, is precisely to say what our line on (what our position on) epistemic matters should be.

But, perhaps as much as anything, what makes the propriety of the Requirement of Deductive Cogency important is the way in which the question as to its propriety engages one of our most critical epistemological problems: the problem of deciding how, as epistemologists, we are to come to terms with our human fallibility. It is, after all, our human fallibility that is to blame for the fact that we cannot expect to tell a story of any significant ambition (that purports to describe how the world is) without having told a story that is most probably false. If Descartes had been right, there would be no call for Christensen and me to dispute the propriety being willing to tell such a story. Descartes thought that, if we just followed his instructions, we would be able to tell an ambitious story about how the world is that comes with a

guarantee that the story is true in every detail. But that is because Descartes thought we are possessed of a cognitive capacity that, if just properly deployed, will infallibly deliver us a significant part of the truth about how our world works.

Of course, we have no such capacity. And it is because we have no such capacity that any story of any ambition about how the world works we might want to tell will contain a great many independent claims of whose truth we have no business being fully confident—and so we cannot but be confident that not all the claims of which the story is comprised are true.³³ It is the question as to how an epistemology is to accommodate the fact that we have no such capacity as Descartes thought we have—how an epistemology is to account for the fact that we are humanly fallible—that Christensen and I are debating.

My own approach to the matter owes a debt to J. L. Austin.³⁴ Austin was also concerned with how an epistemology should accommodate our human fallibility. The worry he was concerned to address is brought out by the following line of reasoning. Suppose I claim to know that there is rosebush in my garden. Even as I make that claim, I must concede that it is imaginable that I am mistaken. It is imaginable that the bush has been removed since I last cast my eyes in the direction of the place I planted it, imaginable (supposing my eyes are now cast that way) that it has been replaced by an ersatz bush that I cannot distinguish from the original at this distance. This being so, it would seem to behoove me to admit that I might be wrong about there being a rosebush in my garden. But once I have admitted this much, it would seem I ought to withdraw my claim to knowledge. I cannot say, “I

³³ This is captured nicely by the idea that it is a condition on your investment of confidence in propositions being warranted that it answers to the axioms of probability. Consider a story that contains even twenty independent propositions in each of which you have good reason to invest a great deal of confidence, but short of the maximum possible. Let us suppose (for simplicity) you have good reason to have a degree of confidence equal to 0.95 in each, and that that the propositions are independent in the sense that your degree of confidence in each, given any conjunction of the rest, is the same as your unconditional degree of confidence. (One way of thinking of this is that you regard each as highly probable, and as the rest as having no evidential bearing on its truth.) It is doubtful that any history of any ambition you might write or read could contain as few as twenty propositions that meet this condition. Yet modest though the story is, you will open your states of confidence to the charge that they violate the axioms of probability unless you have a degree of confidence greater than 0.74 that the story is *false*—and, in particular, that at least one of the twenty propositions is false. Increase the number of propositions, and your confidence that the story is false will have to be greater still.

³⁴ See Austin 1979, esp. p. 98.

know there is a rosebush in my garden, but I might be wrong.” If I know, I *can’t* be wrong.

Austin’s response was to concede that a fair look at our practices reveals that we are perfectly happy to say, “If I know I can’t be wrong.” But, he maintained, a fair look at our practices will also reveal that, we cannot possibly regard “I can’t be wrong”—as it occurs in, “If I know I can’t be wrong”—as something whose falsehood can be established simply by noting that I am humanly fallible. For if it *were* something whose falsehood could be established that way, the question as to whether I may be wrong would never be exciting. But it *is* sometimes exciting. It can be difficult and even, in cases in which important things are at stake, agonizing to decide whether I may be wrong that P—and so cannot claim to know (as perhaps I did) that P. If the fact that I am humanly fallible were sufficient to decide the matter, there would never be a difficulty: there is no question but that I am humanly fallible.³⁵

His view was that we have standards (rough to be sure) for what it takes to sustain a claim to knowledge of a certain sort (where both the content and of the claim, and the context in which it is made, decide what sort it is). These standards are not unreasonably high—they don’t demand, say, that I do something special (more than I’ve already done) in the case at hand, to rule out the possibility that I’ve mistaken an ersatz rosebush for a real one. It’s not that I never make mistakes of this sort. On the contrary, I am *always* capable of making mistakes of this sort. It’s rather that our standards reflect the fact that we happily make and credit knowledge claims—this because we are perfectly entitled to make and credit knowledge claims—even as we fully appreciate that we are always capable of making mistakes. Only when there is a *special reason* to suppose I have made some particular mistake—that is, a reason that is *not* always present when we happily and appropriately make and credit a knowledge claims of the sort I’ve made—can my claim be legitimately challenged on the grounds that I’ve not (yet) done anything special to rule out the possibility that I have made this mistake.

So it is, on Austin’s view, that I can claim to know I’ve a rosebush in my garden, even as I admit that I am humanly fallible and, so, that it is in an important sense possible that I am wrong. When it comes to

³⁵ Nor, one might add, would there be any question as to whether, when (for example) asked breathlessly on the street if I know where the nearest hospital is, I can answer in the affirmative. The obvious answer would be that, of course, I cannot. How could I possibly know where the nearest hospital is when I am a fallible human being and so might be wrong? But that, of course, strikes us as an outrageous way to think about the matter.

deciding whether I might be wrong, for the purposes of deciding whether I know a given claim, the mere fact that I am humanly fallible simply does not settle the matter. Absent any legitimate reason to suppose that I have failed to do enough to sustain my claim to knowledge, it would take a *special reason* to suppose I am mistaken in this case—a reason for supposing I might be wrong that one doesn't *always* have when knowledge claims of the sort I am making are appropriately made and are credited—for it to be appropriate to challenge to my claim on the grounds that I might be wrong. And the mere fact that I am humanly fallible—a fact of which we are always fully aware as we happily and appropriately make and credit claims to knowledge—is not such a reason.³⁶

As I see it, Christensen is appealing to our human fallibility by way of constructing the makings of yet another skeptical trope.³⁷ As I have argued, the story our author has told in her book doesn't count as *being* her story unless she's prepared to stand by everything she's said in the book—and that requires her to be willing to say that everything she's said in the book is true. Christensen is arguing that she cannot in good conscience say such a thing. Why? Because she ought not be willing to say that P if she is aware that it is improbable that P. She surely *is* aware that it is improbable that everything in her book is true. So she ought not be willing to say that everything she's said in her book is true.

It is, to my way of thinking, an argument to the effect that, once we fully appreciate our human fallibility, we must deny that our author—we must deny that any of us—can possibly be warranted in having *any* story of *any* ambition that purports to say how the world is.

My way with this argument has closely followed Austin's way with the one with which he was concerned. I have argued that the

³⁶ I offer an exposition and defense of Austin's response in Kaplan 2006b and Kaplan 2008. I recognize that it may seem odd, in light of my stated commitment to Bayesian decision analysis, that I should be appealing here to Austin's treatment of knowledge. Bayesians have tended to have difficulty seeing what place there is for knowledge in epistemology—how a rational inquirer's ability to fulfill her duties as an inquirer would be in any way compromised were she to attend only to her states of graded confidence, and never concern herself with what, if anything, she knows. (See, in particular, Jeffrey 1968, p. 166.) It may also seem odd that someone committed to Bayesian decision analysis would engage to the extent I have done in Austin-style ordinary language philosophy. Suffice it to say that I think the appearance of oddity can be dispelled. In Kaplan 2000, pp. 300–4, Kaplan 2006a, Kaplan 2009 and Kaplan 2010, I explain how.

³⁷ Although, of course, Christensen does not mean to be doing anything of the sort.

principle, “You ought not say that P in the context of inquiry if you are aware that P is improbable,” is perfectly acceptable, so long as one doesn’t regard “P is improbable”, as it occurs in the principle, as something whose truth can be established simply by noting that P says an awful lot about the world and that we are humanly fallible.³⁸ It is, after all, our practice (and appropriately so) to stand behind the books we write in the full recognition of our human fallibility. The principle holds true only in cases in which you have a *special reason* to regard P as improbable—that is to say, at the very least, a reason that you don’t routinely have for ambitious stories that it is perfectly appropriate for an author to stand behind. But, of course, that we are humanly fallible does not count as such a reason. So, in the only cases in which the principle holds true, our awareness of our human fallibility does nothing at all to suggest that we cannot in good conscience stand behind an ambitious story that purports to say how the world is—and so nothing at all to suggest that we can not say, of such a story, that it is true.

In my view, Austin’s achievement was to see that the recognition that we are humanly fallible doesn’t take the form of an alarm. It, rather, takes the form of background noise to the cries for attention issued by those things that require active consideration as we decide what we know. To be sure, the recognition of our human fallibility shapes our practice. We are open—our human fallibility requires us to be open—to the prospect of finding a special reason to suppose we’ve made a mistake. But that is the *only* way in which the recognition of our human fallibility shapes our practice. Our efforts to decide what we know, while sensitive to special reasons to suppose we are prone to make mistakes in this circumstance or that, quite properly proceed without any further concern for our (general) human fallibility.

What I have been arguing is simply this: that what Austin saw, in our practice of deciding what we know, is present in equal measure in our in our practice of deciding what we ought to be willing to say in the context of inquiry.³⁹

³⁸ Or, as in the case of Christensen’s five improbabilities, by noting that P is probable only if Q is, that Q says an awful lot about he world, and that we are humanly fallible.

³⁹ Versions of this essay were presented at an author-meets-critics session on Christensen 2004 held at the meetings of the Pacific Division of the American Philosophical Association in March 2006, and at the Epistemic Value Conference hosted by the University of Stirling in August 2006. I would like thank the following for conversation and criticism: David Christensen, Gary Ebbs, Ken Gemes, Jesper Kallestrup, Matthew McGrath, Scott Sturgeon, and Joan Weiner.

References

- Austin, J. L. 1979. "Other Minds," in J. L. Austin, *Philosophical Papers*, 3rd. ed. J. O. Urmson and G. J. Warnock (eds), Oxford University Press, pp. 76–116.
- Christensen, David. 2004. *Putting Logic in its Place: Formal Constraints on Rational Belief*. Oxford University Press.
- Hawthorne, James and Bovens, Luc. 1999. "The Preface, the Lottery, and the Logic of Belief", *Mind*, 108: 241–264.
- Jeffrey, Richard C. 1968. "Probable Knowledge," in Imre Lakatos (ed), *The Problem of Inductive Logic*. North-Holland Publishing Company, 166–80; reprinted in Jeffrey 1992, pp. 30–42.
- . 1983. *The Logic of Decision*, 2nd ed. University of Chicago Press.
- . 1992. *Probability and the Art of Judgment*. Cambridge University Press.
- Kaplan, Mark. 1981. "A Bayesian Theory of Rational Acceptance," *The Journal of Philosophy* 78: 305–330.
- . 1996. *Decision Theory as Philosophy*. Cambridge University Press.
- . 2000. "To What Must an Epistemology be True?" *Philosophy and Phenomenological Research* 61: 279–304.
- . 2002. "Decision Theory and Epistemology," Paul Moser (ed), *The Oxford Handbook of Epistemology*. Oxford University Press, pp. 434–62.
- . 2006a. "Deciding What You Know," in Erik J. Olsson (ed), *Knowledge and Inquiry: Essays on the Pragmatism of Isaac Levi*. Cambridge University Press, pp. 225–240.
- . 2006b. "If You Know You Can't Be Wrong," in Stephen Hetherington (ed), *Epistemology Futures*. Oxford University Press, pp. 180–98.
- . 2008. "Austin's Way with Skepticism," in John Greco (ed) *The Oxford Handbook of Skepticism*. Oxford University Press, pp. 348–71.
- . 2009. "Williamson's Casual Approach to Probabilism," in Patrick Greenough and Duncan Pritchard (eds) *Williamson on Knowledge*. Oxford University Press, pp. 122–39.
- . 2010. "John Langshaw Austin," in Sven Bernecker and Duncan Pritchard (eds) *The Routledge Companion to Epistemology*. Routledge, pp. 798–810.
- Kyburg Jr, Henry. 1961. *Probability and the Logic of Rational Belief*. Wesleyan University Press.
- Maher, Patrick. 1993. *Betting on Theories*. Cambridge University Press.
- Makinson, D. C. 1965. "The Paradox of the Preface," *Analysis* 25: 205–7.
- Roorda, Jonathan. 1997. "Fallibilism, Ambivalence, and Belief," *The Journal of Philosophy* 94: 126–55.

- Savage, Leonard J. 1972. *The Foundations of Statistics*, 2nd ed. Dover Publishing.
- Sturgeon, Scott. 2008. "Reason and the Grain of Belief," *Nouûs* 42: 139–65.
- Weintraub, Ruth. 2001. "The Lottery: a Paradox Regained and Resolved," *Synthese* 129: 439–449.
- Williamson, Timothy. 2000. *Knowledge and its Limits*. Oxford University Press.