

Coming to Terms with our Human Fallibility: A Response to Christensen

Mark Kaplan

I. CHRISTENSEN'S CHALLENGE

Of what relevance is logic to a rational inquirer? A standard answer is this. It is part and parcel of being a rational inquirer that one should want one's beliefs to be deductively cogent: that the set of one's beliefs contain all its consequences but contain no contradiction. In his elegant little book, *Putting Logic in its Place*,¹ David Christensen argues that this answer is mistaken: commit yourself to the deductive cogency of your beliefs and you commit yourself to crazy beliefs.

But, for Christensen, the result that logic does not have *this* purchase on our states of belief is not terribly worrying. The sort of belief that deductive cogency has been thought by some to constrain is the sort of belief (Christensen calls it "binary belief") that one either has or doesn't have in a proposition, *tertium non datur*. And Christensen thinks that this is not a sort of belief an epistemology has any call to take seriously. To be sure, we have states of *graded confidence*. And, to be sure, these states of graded confidence are *indirectly* constrained by logic. But that is only because a rational inquirer ought to conform her graded beliefs to the axioms of probability—and those axioms make essential reference to logical properties.² In contrast, the requirement of deductive cogency would have to constrain binary beliefs. But what are they? Once a rational inquirer has determined as best she can how confident she is in a proposition P, there doesn't seem any settling of belief as to P's truth-value left to do. There doesn't seem to be any point to asking--because there doesn't seem to be anything left to be found out by asking--"Yes, but does she *believe* P?"³

Now this not the whole of Christensen's brief. But it is the brief I want to take up. In particular, I want to take up the last bit. I read it as a challenge: a challenge to say how epistemology would be any the poorer were it to (i) give up on the picture of rational inquirers as having states of opinion that answer to the requirement of deductive cogency, and (ii) pursue its concerns strictly in terms of graded confidence. It is a challenge that, Christensen means to convince us, cannot be met.

What I want to assess is how well he has executed that task. But, before I do that, I want to make as clear as possible what the rules of the game are supposed to be: what exactly one is required to do order to count as having succeeded in doing what, according to Christensen, cannot be done.

Clearly, one has to identify a respect in which, even after she's determined how confident she is that P, an inquirer still has some settling of opinion to do. Settling her opinion in this respect has got to be something that is important to an inquirer. It has to make sense that an epistemology would concern itself with settling opinion in this respect. It has to make sense that the propriety of how a person settles her opinions in this respect, and the propriety of the opinions upon which she settles, be legitimate subjects of epistemic appraisal. And, finally, it has to be the case that (in some important

¹ Christensen 2004. All the page references in the text and footnotes are to this book.

² Consider two of the axioms: (1) $\text{prob}(P) = 1$ when P is a tautology; (2) $\text{prob}(P \vee Q) = \text{prob}(P) + \text{prob}(Q)$ when (P & Q) is a contradiction.

³ See, especially, p. 98.

sense) the propriety of how a person has settled her opinions in this respect depends on the satisfaction of the requirement of deductive cogency.

But does one have to come up with an account, of what settling opinion in this respect comes to, that will answer to *everything*, or even *most of* what, we are ordinarily inclined to say when we talk about beliefs? It seems to me that the answer has to be “no”. We can surely demand that, whatever an account says that settling opinion in this respect comes to, it be something we have used our talk of beliefs to talk about. But it may be that our ordinary talk of belief is hopelessly confused, rent by our commitment to competing psychologies. If so, there can be no coherent account that is entirely faithful to our ordinary use of “belief”.⁴ And yet the challenge will still be live. For we can still ask, “Is there *any* use we have been making of the expression ‘belief’ that points to an interesting respect in which a rational agent, who has already determined how confident she is that P, may yet still not have settled her opinion with respect to P?” If the answer is “Yes,” then the story of how one ought rationally settle one’s opinions in this respect deserves a place in our account of a rational inquirer. And should it turn out that part of that story is that (in this respect) a rational inquirer’s opinions are subject to the requirement of deductive cogency, it will have then turned out that logic has a greater claim on a rational agent than Christensen thinks.

II. A BIFURCATION ACCOUNT

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.

Of course, in different circumstances—in circumstances in which I was engaged in a performance designed to amuse, or shock, my audience—I might very well do what I just said I wouldn’t do.⁵ The reason I wouldn’t do it on the occasion I am imagining—the reason it would be outrageous for me to do it on such an occasion--is that, on that occasion, I am saying what I am saying, and doing what I am doing, in the context of inquiry. And while, in that context, I might be very happy to amuse you or shock you, I am constrained by what means I can do it. I open my position to criticism if I have sought to amuse you or shock you by saying something or doing something that conflicts with what is reasonable to say or do in the context of inquiry—what is reasonable to say and do when one’s aim is (roughly speaking) just to tell (as best one can) the truth as it pertains to the matters at hand. 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.

⁴ I have argued that our ordinary talk of belief *is* incoherent in just this way. See Kaplan 1996, p. 154; 2002, p. 450.

⁵ Indeed, I can easily imagine John Cleese in a Monty Python sketch doing exactly what I said I wouldn’t do; he did things not so dissimilar.

But, actually, a stronger principle is warranted. Suppose that you did something different: you pointed to a set of things I've said, and observed that this set logically implies P. And then you asked, "But what *reason* have you for saying that P?" No charge of inconsistency being leveled here. Yet it is a legitimate challenge. 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 argue for 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, they logically imply P. But I'm not *saying* that P. And, only if I were *saying* that P, would I need a *reason* for saying 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 P is something that (for example) there is no good reason to say. 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) *deductively cogent*: i.e.,

- (i) the set should contain everything that set logically implies; and
- (ii) it should contain no contradiction.

Why do I think it is of interest, for present purposes, that the set of things you are willing to say in the context of inquiry is thus constrained by the Requirement of Deductive Cogency? Because, I want to suggest, it is in precisely this respect in which your opinions are subject to the demands of deductive cogency. And I want to suggest that, once we recognize as much, we have all we need to meet Christensen's challenge.

Recall what it takes to meet the challenge.

First, one has to identify a respect in which, even after she's determined how confident she is that P, an inquirer still has some settling of opinion to do. This constraint is satisfied: even after you have determined as best you can how confident you are at that P, there is still the matter of deciding what, if anything, you are willing to say, in the context of inquiry, as to whether it is the case that P.

Second, settling opinion in this respect has to be something we have used at least some of our talk of beliefs to talk about. This constraint is satisfied too. Consider that fact that, in the context of inquiry, we treat talk of whether we want to say that P, and whether we believe that P, as entirely interchangeable. Consider, as well, the fact that we find it natural to suppose that the preface paradox--which would seem, in the first instance, to be concerned with the extent to which what we are willing to say in the body of a book constrains what we are willing to say in the book's preface--expresses a concern as to how our *beliefs* are constrained.⁶

Third, settling your opinion in this respect has got to be something that is important to us as inquirers: it has to make sense that an epistemology would concern itself with settling opinion in this respect—that the propriety of how a person settles her

⁶ The Preface Paradox was first presented in Makinson 1965. The paradox arises from thought is 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.

opinions in this respect, and the propriety of the opinions upon which she settles, be legitimate subjects of epistemic appraisal. This third constraint is satisfied as well. It matters a great deal to us as inquirers what it is we, and others, are willing to say, and ought to be willing to say, in the context of inquiry.⁷ 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 in them, testifies to our preoccupation with how, in this respect, we should settle our opinions.

Finally, how we in this respect settle our opinions must be constrained in some important way by the requirement of deductive cogency. And, as I have just argued, this condition is satisfied too.

The fact that, on the view I am proposing, how confident you are that P doesn't immediately settle, in the relevant respect, your opinion as to whether P is the case, makes me (what Christensen calls (27)) a bifurcation theorist: how confident you are that P, and whether you are willing to say that P in the context of inquiry, are two different matters.

What I like about being a bifurcation theorist is that it gives me the freedom to accommodate aspects of our intellectual life that simply don't fit very well in a view that would make the propriety of being willing to say that P, in the context of inquiry, a byproduct of the propriety of being, with respect to some fairly uniform threshold, sufficiently confident that P.⁸ On the latter sort of view, we might do a perfectly adequate job of determining what we should say in the context of inquiry in a completely piecemeal way: by simply examining each proposition one at a time and determining whether our confidence in its truth exceeds the threshold. To be sure, there will be a consistency constraint and an epistemic constraint at work here: we will want to check that our investments of confidence conform to the probability calculus, and we will want to check that our investments of confidence are appropriate to our evidential situation. But there is no call, on such a view—no place on such a view—for us to evaluate propositions and hypotheses *qua* things we might be willing to say.⁹ There is no call—no way—to consider how they hang together *as things to say*. What we are willing to say in the context of inquiry is nothing more than an epiphenomenon of the shape our confidence assignment takes.

This seems to me just plain false to the way we actually decide what we should be willing to say in the context of inquiry. There, it seems to me, we are looking, not for some individually probable things to say, but for a deductively cogent story to tell about how the world is.¹⁰ And we evaluate hypotheses 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

⁷ To be sure, what things you are willing to say in the context of inquiry is of nowhere near as great consequence as how confident you are, for each of those things, that it is true. The former matters only in the context of inquiry. The latter matters in every context, practical or intellectual, to which any of the hypotheses in question is germane.

⁸ An example of such a view: one is justified in being willing to say that P in the context of inquiry if and only if one is justified in believing P (where belief is just a state of graded confidence above a fairly uniform threshold less than 1).

⁹ Or, more properly, to say *to be true*.

¹⁰ This way of putting it is due to Jonathan Roorda. See Roorda 1997, pp. 148-9. Christensen quotes (73) the relevant passage.

things that go into it, are worthy of our confidence. But 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—we are looking for a story that gives us a line on things. So, as we assess ways we might alter our story, we weigh competing goods: the good of saying something true, and the good of saying something contentful.

Of course, the sort of deductively cogent story that I've described us as looking to tell in the context of inquiry is not the sort of thing in whose truth we will have any business being confident. A big, ambitious, story brings with it a very great risk of error. But a bifurcation view—a view that recognizes being confident in P as being one thing, and being willing to say that P in the context of inquiry another—can make sense of our willingness to tell such a story in the context of inquiry. What makes it rational for us to be 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.¹¹

As I say, this seems awfully nice to me. But by Christensen's lights, the very things that I find attractive about being a bifurcation theorist—the freedom such a theory affords me to accommodate the Requirement of Deductive Cogency, the freedom such a theory affords me to make sense of our willingness to tell a big story in the context of inquiry even though we recognize how unlikely it is to be true—make a bifurcation theory intolerable.

III. DEDUCTIVE COGENCY REVISITED

Let me begin with the Requirement of Deductive Cogency. The consideration I have offered in its support comes to this: we cannot account for the force *reductio* arguments exert upon us unless we endorse this requirement. Christensen denies this. He maintains (79-96) that the force of *reductio* arguments can be fully explained by a theory that sees a person's willingness to say that P in the context of inquiry as nothing more than an expression of a degree of confidence in P above a threshold¹²--that is, by a theory on which deductive cogency has no purchase on what a person ought to be willing to say in the context of inquiry.

He notes that on such a theory, a *reductio* of a small number of hypotheses can constitute a very serious criticism, showing that there is no way consistent with a probabilistic constraint on confidence to invest enough confidence in each of the hypotheses involved so that they will all exceed the confidence threshold. Only *reductios* that require very large numbers of hypotheses lack critical bite. But, maintains Christensen, such *reductios* lack force anyway.

¹¹ In Kaplan 1996, pp. 123-4, I offer a sketch of how our ordinary standards for evaluating such stories reflect the fact that we make just such trades. I then go on to say why this trade-making does not require that content actually be something that is subject to measure.

¹² Theories of this sort admit of two varieties: ones that *define* a belief in P as nothing more than a degree of confidence in P above a threshold, and ones that (however—or, indeed, whether or not—they define belief) hold that belief in P is *warranted* if and only if the bearer of the belief has a warranted degree of confidence in P above a threshold.

I agree with this much. The confidence threshold accounts under discussion can explain the force of relatively small *reductios*—and that they cannot do similarly for ones with large numbers of premises. I also agree (as I have already said) that there is no reason why the mere fact that a large set of hypotheses admits of a *reductio* should disincite us to invest each of its premises with a very high degree of confidence. Witness the lottery paradox.¹³ All I deny is that, so long as it has sufficiently many premises that each can rationally receive a degree of confidence above the threshold, we are free to dismiss a *reductio* meant to criticize our being willing to say, in the context of inquiry, the things that serve as the argument's premises.

To explain why, I have appealed to the following case.¹⁴ Suppose that you charge me with producing a chronology of the events that preceded a serious accident. 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 hypotheses 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 hypotheses. The idea that only *reductios* of a relatively small number of hypotheses has critical bite, when it comes to what can rationally be said in the context of inquiry, simply won't wash.

Christensen's response, as I read it,¹⁵ is that the case as I described it is not detailed enough to permit us to determine whether my chronology is really open to criticism. 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 that the chronology is defective. (89-90)

In this, it seems to me Christensen is just plain wrong. My production of the chronology opens my position to criticism no matter *why* I produced it and why I am as confident as I am in each of its elements. It opens my position to criticism even if, as is

¹³ The Lottery Paradox was first stated in Kyburg 1961, p. 463. Imagine that the inquirer is faced with an million-ticket that she is justified in being certain is fair and will produce exactly one winning ticket. She will be justified in having a degree of confidence equal to 0.999999 in each of the first n members of the following set of propositions (and a degree of confidence equal to 1 in the last): {ticket 1 will lose, ticket 2 will lose,..., ticket 1,000,000 will lose, not all of tickets 1 through 1,000,000 will lose}. But this set of propositions represents a set of things she cannot be willing to say unless she violates the Requirement of Deductive Cogency. For the set of propositions is inconsistent.

¹⁴ Kaplan 2002, pp. 459-60, n 20. I reproduce my description of the case *verbatim*.

¹⁵ I say, "as I read it", because he doesn't actually describe the case as I just did—and as I did originally. He describes a case in which you ask me to come up with a many-part chronology, and the chronology above is what I produce (89). In the case as I described it, you simply ask me to come up with a chronology. The decision to place it many parts is my own. In so recasting the case, I think he misses its real import. See below.

perfectly imaginable, I am entirely rational to be as confident as I am: I find each of the hypotheses 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. This one does not.

Deny this, and we license, as a sound strategy to insulate what one wants to say from criticism, the decomposition of each individual thing one wants to say in the context of inquiry into a set of hypotheses each of which is as modest as is compatible with the set's conjunction being equivalent to the original hypothesis thus decomposed.¹⁶ "Don't say that cigarette smoking causes cancer," it will say on the license we issue. "Someone might be able to show that you have no business being sufficiently confident to say that it is so. Instead, say—or, more conveniently, declare your willingness to say--each of the following million things:¹⁷ 'Either cigarette smoking causes cancer or ticket #1 will lose in the lottery'; 'Either cigarette smoking causes cancer, or ticket #2 will lose in the lottery';... 'Either cigarette smoking causes cancer, or ticket #1,000,000 will lose in the lottery.' No matter what, no one will shake your confidence in anything you'll be saying now." We license a pattern of behavior that, if we and all our colleagues in all our universities were to adopt it and succeed in teaching it to our students, would pervasively affect the conduct of inquiry in a way that would render it unrecognizable and (at least, by the lights of our current sensibilities) pointless.

IV. THE PREFACE

Of course, to endorse the Requirement of Deductive Cogency is to claim that an inquirer 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. Christensen thinks it a mad result of my view that it should require her to be willing to say something so improbable, so outrageously immodest.

But is it really? Christensen admits (35) that, despite the fact that (as she must concede) it is overwhelmingly probable that our author has made some mistake in the body of her book, we would find it strange if she flatly said in the preface, "This book contains errors." He suggests (43) that this may be because (as some writers have claimed) we understand unqualified assertions to express knowledge, not merely belief. And so, he suggests, the awkwardness we see in her saying flatly that her book contains 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 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 should 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?" That's because our expectation is that an author will flatly say that there is an error in her book only when she has a special reason

¹⁶ It is the illegitimacy of doing precisely this that I meant the chronology case to highlight, and that Christensen's recasting of the case (see the previous footnote) obscures.

¹⁷ I am supposing here that you are in the circumstances imagined in footnote 15: you are confronted by what you are certain is a fair million-ticket lottery which will produce one winning ticket.

to suppose the book contains one—that is to say, a reason one doesn't *always* have with respect to a book of any significant ambition. 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.¹⁸ As things stand, she has no special reason—no reason one doesn't *always* have with respect to a book of any significant ambition—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 said--it does

My thought is that, similarly, our author has no call to shrink from saying that everything in the book is true, once she has read it through it and checked everything she has said therein.¹⁹ She has no call to shrink from saying so unless she has some special reason—a reason one doesn't *always* have with respect to a book of any significant ambition—to suppose it contains a mistake. She should stand behind everything she's said until someone has come up with a special reason to suppose there is a mistake. And to stand behind everything she said just *is* to be willing to say, until someone provides a special reason to suppose otherwise, that everything she's 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.²⁰

¹⁸ 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.

¹⁹ She has, of course, no particular obligation actually to say it in the preface of her book.

²⁰ There are a number of worries one might have about this last claim.

First worry: it is too sweeping a claim. What of a book, or set of things one might say, that doesn't seem to take the shape of a story, but is rather just meant to list a set of facts—say a phone book, or set of reports of what horse won which race? [It may be no accident that Christensen adverts to books (35n) and sets of things one might say (92n) of just this sort in criticizing the propriety of the Requirement of Deductive Cogency.] But I don't see a relevant difference. The author is still telling a story as to how the world is (what person has what phone number, what horse won what race). And it isn't her story unless she stands behind it, unless she's willing to say that everything in the story is true.

Second worry: the claim is false. Christensen seems to be of this opinion. 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. But that is perhaps 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. It is a view on which, as we saw, it is a matter of indifference whether a story is even consistent, so long as the story has been told in sufficiently small pieces that one needs to deploy of a large number of those pieces to generate a *reductio*. It is a view on which, as the chronology case shows, it is difficult to make sense of why our author would care to write a book at all.

Third worry: for all that, were our author unwilling to say that the story she tells in her book is true, but only willing to say that it is approximately true, we wouldn't have to interpret her as being

What of Christensen's charge that 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? It seems to me unfounded. After all, our author is still free to engage in the graceful practice of exculpating those whose aid help 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 (because it is the only thing that would violate the Requirement of Deductive Cogency) 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.²¹ 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.²²

unwilling to stand by that story—that is to say, unwilling to take full responsibility for the consequences of her story. But my sense is that we would. Suppose she is willing to say nothing stronger than, "The story I've told in my book is approximately true." Now suppose we show her that a minor consequence of her story is false. It doesn't sound (at least, to my ear) outrageous for her to say in response, "I grant your point, but that does not materially affect the propriety of what I've said. I never said the story is true in its every detail. I said the story is *approximately* true. And what you've shown me doesn't do anything to show that my story is *not* approximately true." That is, it does not seem outrageous for her to dismiss, as no criticism at all of her having said her story is approximately true, that she has been forced to say something (the point she's granted) that, together with the story she's told, is open to a *reductio*. But then she's taking just the sort of limited responsibility for the consequences of her story that, as the chronology case shows, makes it hard to make sense of her book as being meant to tell a story. It is hard to see in what sense a person who is writing things down, with indifference to whether the set of things she writes is consistent, really counts as even *trying* to tell a story.

²¹ 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, "I never said it did!?" 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 this way: "The Coroner's Office takes responsibility for the dead, *if there are any.*"?

²² One reason Christensen thinks that the Requirement of Deductive Cogency requires our author to be unduly immodest is that he thinks (46ff.) that, once she adopts a willingness to say that everything she's said in her book is true, she is committed to making unduly immodest comparisons between her book and books written by others. He thinks she is thereby committed to regarding her own book as being unique in having managed to emerge error-free, and thus superior even to books authored by (what she herself recognizes to be) scholars more meticulous than she. She is committed to being willing to write, in her proposals to granting agencies, that she's written the first error-free book. Supposing that there is a prize on offer to the author of the first error-free book, our author is committed to be willing to say in the context of inquiry that she is going to win that prize.

But Christensen is mistaken. True enough, the Requirement of Deductive Cogency requires our author to be willing to say, of her own book, that it contains no errors. And, true enough, the Requirement of Deductive Cogency does nothing to keep her from saying, of every *other* book, that it contains errors. But neither does the Requirement of Deductive Cogency do anything to *require* her to be willing to say, of every other book, that it contains errors. That is, there is nothing in the Requirement of Deductive Cogency that requires her to be willing to say that hers is the only book ever written, the first book ever written, that is error-free.

Indeed, it seems to me that 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 *always* have with respect to a book

But how exactly 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 being 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.”

But 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 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 *always* present in the

of any significant ambition, 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.

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 checked as she has her own. (She has read her own book sentence by sentence and checked that each sentence expresses something she is willing to say.) She will be willing to say her own book 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 checked as she has her own.

The upshot, then, is that the Requirement of Deductive Cogency does not impose upon our author the undue immodesty Christensen thinks it does: he thinks the Requirement forces our author to be willing to say that hers is the first and only error-free book. To do that, the Requirement would have to force our author *both* to be willing to say that her book is error-free *and* to be willing to say, of every other book, that it contains errors. But, as I have pointed out, the latter is not anything the Requirement of Deductive Cogency demands of our author. Indeed, I have argued, she has no business being willing to say, of a book she hasn’t read, that it contains errors.

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--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 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 overwhelmingly 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 everything I've said in this book. Ask me if everything I've said in the book is true, and my answer is, 'Yes, it is true.' Naturally, I recognize that mine is a book of significant ambition. And I recognize that it is in the nature of a book of significant ambition that it most probably contains some error or other. But I have no special reason to suppose it contains any errors. On the contrary, I have checked it sentence by sentence for errors and have found none. That is why, though I recognize the extremely high probability that it contains some error or other, I stand by the story I have told in the pages to follow and am prepared to say that everything I say there is true. Of course, I am prepared alter my stance as soon as I find any special reason to suppose there is a mistake."

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 has 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 is 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?²³

I do not mean to suggest that there is *nothing* our author could say by way of providing a suitable commentary that will defeat the expectation that would be fostered by her flatly saying, “This book contains errors.” On the contrary, it seems to me she could say something like this: “I am not aware of any errors in this book. I have checked quite thoroughly and haven’t found any. 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.”²⁴

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 thought reflected in the commentary reproduced at the end of the last paragraph—in 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.²⁵

V. COMING TO TERMS WITH OUR HUMAN FALLIBILITY

One may be tempted to think that all this back-and-forth, about whether there really is any place in our account of a rational inquirer for a sort of settled opinion that cannot

²³ I should note that it is via an imagined dialogue with an author who *denies* that his book contains errors, that Christensen attempts to bring out the absurdity of an author’s being willing to say that her book is error-free.

²⁴ I thank Scott Sturgeon for suggesting as much.

²⁵ In a longer version of this essay, I address a further complaint Christensen raises against the view I am defending here: that it would have our author behave in crazy ways.

understood in terms of graded confidence, is nothing more than a dispute over how to work out some minor wrinkle in a rather formal little outpost of epistemology. But I don't see it that way.

After all, what is under dispute is how best to understand a practice that, at least for some of us, is *very* important—and very central—to our ordinary lives: the practice of determining what we should say to be so in the context of inquiry. A particular object of my interest is how to make sense of the fact that this practice appears to take the form of seeking an ambitious story (as to how the world is) to tell in the context of inquiry. The question I am interested in is how we can be willing, in good conscience, to tell such a story—thereby committing ourselves to saying that it is true—even though we regard it as highly probable that it is false.

Christensen thinks we *can't*. And so he tries to spin an account, of what it is we are doing when we tell a story in the context of inquiry (say, by writing a scholarly book), on which we *needn't*. What I have been arguing (by way of playing my part in the back and forth about whether there is a place for some opinion, that is not to be understood in terms of graded confidence, in an account of a rational inquirer) is that Christensen's account won't work. His account requires us to give up the Requirement of Deductive Cogency. But if we really did that, it would alter our practice (of seeking a story, about how the world is, to tell in the context of inquiry) in a way that would make that practice unrecognizable and (at least, by the lights of our present sensibilities) pointless.

Our struggle, Christensen's and mine, is (as I see it) a struggle over how 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 nothing over which to struggle. He thought we could tell—he promised to show us how we could tell—an ambitious story, about how the world is, that would come with a guarantee that the story was 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. It is the question as to how an epistemology is to accommodate the fact that we have no such capacity—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.²⁶ 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 a 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 know there is a rosebush in my garden, but I might be wrong." If I know, I *can't* be wrong.

²⁶ See Austin 1979, esp. p. 98.

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.²⁷

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—and think ourselves 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 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 made—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.²⁸

As I see it, Christensen is appealing to our human fallibility by way of constructing yet another skeptical trope. 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

²⁷ 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.

²⁸ I offer an exposition and defense of Austin's response in Kaplan 2006b and Kaplan forthcoming. Of course, it might be wondered how an advocate of Bayesianism can legitimately appeal to Austin, the avatar of ordinary language philosophy. For the answer, see Kaplan 2000, pp. 300-4.

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 in good conscience claim to have 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 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 we are humanly fallible. It is, after all, our practice 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—a reason you don't have for *every* story of any ambition. 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—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.²⁹

Indiana University

²⁹ This is a revised version of an essay written for an author-meets-critics session on Christensen 2004 held at the meetings of the Pacific Division of the American Philosophical Association in March 2006. That essay was, in turn, an outgrowth of a presentation I made, at the invitation of Scott Sturgeon, to his seminar at the University of Michigan in February of 2005. I have benefited from Joan Weiner's comments, from David Christensen's response to my contribution to the APA session (and our long, very useful, conversation that followed the session), and from a series of profoundly helpful discussions--over the course of more than five years--with Scott Sturgeon.

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*. Cambridge University Press.
- Foley, Richard. 1993. *Working Without a Net: A Study of Egocentric Epistemology*. Oxford University Press.
- Hawthorne, James, and Bovens, Luc. 1999. "The Preface, the Lottery, and the Logic of Belief", *Mind*, 108: 241-264.
- Hunter, Daniel. 1996. "On the Relation between Categorical and Probabilistic Belief," *Noûs* 30: 75-98.
- 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," in 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.
- . Forthcoming. "Austin's Way with Skepticism," in John Greco (ed) *The Oxford Handbook of Skepticism*. Oxford University Press.
- Kyburg, Henry, Jr. 1961. *Probability and the Logic of Rational Belief*. Wesleyan 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.
- Sturgeon, Scott. Forthcoming. "Reason and the Grain of Belief," *Noûs*.
- Weintraub, Ruth. 2001. "The Lottery: a Paradox Regained and Resolved," *Synthese* 129: 439-449.