

Closure, Defeasibility and Conclusive Reasons

Claudio de Almeida

Received: 2 October 2007 / Accepted: 26 November 2007 / Published online: 28 December 2007
© Springer Science + Business Media B.V. 2007

Abstract It is argued, on the basis of new counterexamples, that neither knowledge nor epistemic justification (or “epistemic rationality”) can reasonably be thought to be closed under logical implication. The argument includes an attempt to reconcile the fundamental intuitions of the opposing parties in the debate.

Keywords Epistemic closure · Defeasibility · Conclusive reasons · Knowledge · Epistemic justification · Epistemic rationality

In what follows, I argue against every tenable closure claim for either knowledge or justification. This case against epistemic closure is based on new counterexamples and is accompanied by an attempt to accommodate the deepest opposing intuitions on both sides of the closure divide.

1 Elitist Closure and Paradox: A Non-Dretskean Refutation

Although I will be arguing against *every* appealing epistemic closure claim, it will be useful to have some prime instances of our target before our eyes at the outset.¹ So,

¹I confess that I find it downright uncomfortable to speak, as most in the literature have done, of a *principle* that is not an explicit directive. So, I’ll take the liberty to speak of closure *claims* or *theses*.

C. de Almeida (✉)

PUCRS (Brazil), Rua Gonçalves Dias 319, ap. 901, Porto Alegre, RS 90130-061, Brazil
e-mail: cdeal@terra.com.br

consider the following theses, some of the most popular in the contemporary literature on the issue²:

- (a) “[R]oughly, if you know p and you know that p implies q , then you know q .” (Roush 2005; also Pritchard 2005, Stine 1976, Dretske 1970).
- (b) “If e is sufficient evidence for S to know q (on the basis of e) and S knows that q entails not- h , then e is sufficient evidence for S to know not- h (on the basis of e).” (Cohen 1988; also Dretske 1970, 2005)³.
- (c) “If S is justified in believing P_1, \dots, P_n , and knows that P_1, \dots, P_n jointly imply Q , then S is justified in believing Q ” (Olin 2005; also Wright 2002).
- (d) “If S is justified in believing P , and knows that P implies Q , then S is justified in believing Q ” (Olin 2005; also Canary and Odegard 1989).

We shall discuss other claims in the closure family shortly. But these are a good starting point, for they prompt an important anti-closure point, one that should have been made a long time ago but got lost in the conceptual fog that covers much of the ground between philosophical logic and epistemology.

Insofar as known implication is a major factor in epistemic closure, paradoxes involving logical implication (that is, paradoxes involving valid deductive reasoning) seem to offer us a knockdown case against justified belief in closure claims. Consider the most stunning case I can give you: Russell’s paradox.⁴ Recall the crucial steps in Russell’s paradoxical argument. We *know* — at the very least, we clearly seem to have all it takes to know — that some classes are members of themselves. Only the despair induced by the awareness of a paradox, we can now see from our vantage point, could possibly have led otherwise-level-headed, even venerable people to denounce such a platitude. For instance, the class of all non-philosophers is not itself a philosopher, and so — *common sense alone tells us* — it is a member of itself, being itself one of the items of which it is supposed to be the collection. Likewise, we *know* — at the very least, we clearly seem to have all it takes to know — that some classes are not members of themselves. For instance, the class of all humans is not itself human, and so it is not a member of itself. We can assert these things and not feel at all queasy, fearing we might be talking nonsense. Further, we can, of course, soundly collect all the classes which are members of themselves. The collection will properly be described as being the class of all classes which are members of themselves. Again: We *know* that there can be classes of

² The authors associated with the following theses are only some of those who have *discussed* these theses, occasionally remaining non-committal about those with which they are associated in what follows. The reference is not meant to indicate endorsement of a given thesis by a given author. Also, the association is only a *good approximation* to the author’s actual choice of words. Some of the theses are not found verbatim in these references. The author whose name appears first is the one who is quoted.

³ In his 2005 paper, Dretske writes: “[I]f S knows that P is true and knows that P implies Q , then, *evidentially speaking*, this is enough for S to know that Q is true” (emphasis added). Internal evidence suggests that he should be associated with this claim discussed in Stewart Cohen (1988) — even though he has also subscribed to claim (e), below, in his 2005 reply to John Hawthorne.

⁴ Russell’s paradox is the most stunning of all the paradoxes involving deductive reasoning that I know. To my mind, the other logical paradoxes give us more wiggle room to reject (or suspend judgment on) one of the premises. But I’m giving you a recipe here: Take your favorite paradox involving valid deductive reasoning and it should look to you like a reason for withholding belief in epistemic closure claims.

classes, and we also *know* — at the very least, we clearly seem to have all it takes to know — that some classes can be members of themselves. No problem there, *not even a hint of a problem* from an intuitive point of view. The light of reason shines brightly as you make those assumptions. And, so, there is nothing whatsoever in sight to stop you from thinking that you can, likewise, collect all the classes which *are not* members of themselves and properly describe the resulting collection as being the class of all classes which are not members of themselves. But, as we also know, at this juncture, the genius of Russell noticed that you must also accept that, if the class of classes which are not members of themselves *is* a member of itself, it cannot be, and, if it *is not*, it must be. And, it doesn't take genius to notice that you cannot reasonably go on to draw, from those seemingly irrefutable assumptions, the conclusion logically implied by them that the class of classes which are not members of themselves *both is and is not* a member of itself, just because a contradiction that is recognized as such is as clearly a case of a *rationaly unbelievable and unknowable* proposition as any we have known.⁵ (If this is not rationally unbelievable and unknowable, we just don't know what is!)

All of this can reasonably be accepted without any qualms, *unless*, that is, you are justified in thinking that at least one of those epistemic closure claims are true. If you are, you might understandably feel compelled — and rationally so — to reject some of those paradoxical assumptions. In that case, however, you had better be ready to explain why you can more easily (and reasonably) reject one of those assumptions than your favorite closure claim. But, please, do notice what, exactly, you are up against. I'll mention it briefly.

For a century now, it's been hammered into our heads that a way out of Russell's paradox will take the form of an explanation of why at least one of the paradoxical assumptions is a wolf in sheep's clothing. Specifically, it's been claimed — but never established beyond a reasonable doubt, *not even close* — that there is good, undefeated reason for thinking that at least one of those paradoxical assumptions is *either false or nonsensical*. Now, of course, if you saw anything that passed for a proof of that claim by your standards, you most definitely would not be justified in believing the seemingly false/nonsensical assumption. But, as you also certainly know, the promise of any such proof beyond a reasonable doubt has gone unfulfilled for about a hundred years now. You don't have to take my word for it. I can give you expert testimony to that effect from a philosopher whose expertise is, by any acceptable standards, beyond a reasonable doubt, a philosopher whose conviction in that regard is so firm that he is willing to stake his enviable reputation on it: Graham Priest.

As you may have already noticed, although Priest and like-minded philosophers have put forward what the vast majority of us regards as an outrageous claim — the claim that some contradictions are true (in addition to being false) — no reasonable critic would be so sanguine as to claim that Priest's argument for that claim, relying crucially, though not exclusively, on Russell's paradox, is a *non-starter*.⁶ Priest's argument is obviously

⁵ It is, of course, awkward to speak of a "knowable proposition", and of "knowing a proposition", but this is well-established philosophical clumsiness that allows us to avoid painful circumlocution. I follow standard usage.

⁶ My main reference here is Priest (1986), a philosophically to-the-point case for dialetheism. For the latest version of his case, with replies to criticism and all the nice formal details, see Priest (2006).

not laughable. It is not laughable because we can all easily follow him half-way into it, stopping just short of accepting its outrageous conclusion.

So, there you have it: In capitalizing on Priest's premises — on Priest's reputation and bold argument — my case becomes instantly appealing to all of those who think that Priest's challenge *is* a well-founded one *but* must ultimately fail. Priest's case for rational belief in a contradiction must fail because there is this non-paradoxical way out: Let go of your closure assumption.⁷ From an epistemological point of view, Priest is an extreme case of justification-closure dogmatism. Since the Russellian contradiction is rationally supported by the Russellian argument, believing it, he claims, is rational (justified) for you, even though you see that it is a falsehood. However, because the Russellian argument is perceived as valid with true premises, you are expected to understand that that falsehood is also a truth. So, you're not required by Priest to accept that it is rational for you to believe what you see is *only false*, but only to accept that, in the Russellian case, it is rational for you to believe what is *both true and false*.

Now, you may be thinking: Is it really accurate to describe Priest as “an extreme case of justification-closure dogmatism”? If, like Priest and me, you are not pre-theoretically inclined to question the *validity* of the Russellian argument or the truth of its premises, and yet, like me and unlike Priest, you don't allow for true falsehoods, you would now seem to be in a position rationally to think that this Priestian argument is *unsound, but valid with true premises*:

Russell's paradoxical premises are true.

Russell's paradoxical argument is valid.

Therefore, Russell's paradoxical argument is sound.

(Therefore, its self-contradictory conclusion is true.)

But that, of course, is absurd! The very meaning of “soundness” is compromised here! Haven't we gone wrong in assuming, with Priest, that the Russellian premises are true, or that the Russellian argument is valid, without “going Priestian” all the way?

The beauty of the situation for the closure denier is that he can safely shrug his shoulders in response to that question. No matter what, in view of the conceptual quandary into which we find ourselves, *unless the Priestian way is to be seen as a sorry blunder*, epistemic closure will have to go! Even if you assume that closure denial is as repugnant as the other three alternatives, the three horns of our trilemma — falsehood of the premises, invalidity, and the Priestian line (dialetheism in cahoots with, perhaps ultimately motivated by, justification-closure) — unless you can reasonably exclude all but one of the alternatives, you had better withhold judgment about any of those choices — *including epistemic closure*. That is the case against closure from Russell's paradox *as sponsored by Priest's reaction to it*. So, Priest's dialetheism *can*, after all, fairly be described by us as “an extreme case of epistemic

⁷ Granted, some will feel inclined to hold that this anti-closure move *is* paradoxical. I will give you reason to reconsider this optimistic view of the intuitive content of closure claims in Sect. 3 below.

closure dogmatism”.⁸ And, if you are not a Priest follower, or one who can offer a compelling case for the alternatives, you shouldn’t be a closure supporter either.⁹

So, yes, ultimately, we will have to admit that we cannot reasonably believe that Russell’s paradoxical premises are true — and the momentary temptation to speak paradoxically of a Priestian argument as “unsound but valid with true premises” vanishes. But the important point here is that we should be withholding belief in closure claims in the same breath, for epistemic closure claims are certainly no more compelling than Russell’s premises.¹⁰

Let me emphasize the point being made here. It is not open to you to argue as follows: “Claudio, if you don’t follow Priest all the way, if you see reason to withhold belief in the Russellian premises, you don’t have a case against closure, because closure is (in essence) the claim that, if you *are* rational in believing the premises, then you are rational in believing the conclusion implied by those premises (and, if you’re not rational in believing the premises, all bets are off). But, if, while holding the argument valid, you acknowledge that those premises cannot be rationally believed, closure is not refuted by Russell’s paradox.” You shouldn’t argue like that — precisely because the reason adduced here for thinking that the Russellian premises cannot rationally be believed is the apparent fact that dialetheism, the other closure-preserving move open to us, cannot be either dismissed or embraced easily (which is why we were momentarily attracted to the absurd view that the Priestian argument seemed unsound but valid with true premises). *And there seems to be no other good reason for withholding belief in the Russellian premises!*¹¹ So, that objection would be question-begging. The objection would assume epistemic closure in order to argue, by modus tollens, against one’s knowledge of/justified belief in the paradoxical premises. But it loses its grip on us once closure, itself, is held under suspicion.

For agents who know about logical implication, as envisaged in claims (a) – (d), the argument from Russell’s paradox yields the conclusion that, while ignoring their own tacit acceptance of epistemic closure, they should understand that they face a dilemma (if we disregard the invalidity option): They will either *disbelieve* (reject) the paradoxical premises (or believe that, appearances notwithstanding, they are meaningless) — but it takes Herculean strength to do that — or else *believe* the self-contradictory conclusion (with Priest and the dialetheists) — but it takes Herculean strength to do that. If, on the other hand, in addition to knowing about logical implication, they are not (tacitly or otherwise) committed to the view that epistemic

⁸ In his 1986 paper, Priest explicitly suggests that Russell’s paradoxical argument is sound and that the argument is his basis for saying that “I, for example, believe that the Russell set is both a member of itself and not a member of itself”.

⁹ “Closure denial” is something of a misnomer here. My case from Russell’s paradox is a case for justified *suspension of judgment*. It is, however, tempting to avoid deviating from the established terminology of the debate as much as possible. This is still a case *against closure*, though not exactly a case for closure *denial*. It will, however, turn into a case for closure *denial* in Sects. 5 and 6 below.

¹⁰ I thank Anthony Brueckner, Stewart Cohen, Patrick Greenough, and Doris Olin for discussion leading to this point.

¹¹ Again, for the fate of those who have argued against acceptance of the Russellian premises, I’ll refer you to Priest (2006).

status is invariably transmitted by it, they will not be under pressure to choose between either of those horrendous alternatives.¹²

2 Populist Closure and Paradox: The Refutation Continued

We're beginning to overcome the ever-present fascination with *known implication* in the debate over closure. If we focus on *known* implication, as Fred Dretske originally did, we may make the case against closure more dramatic, or we may, conversely, give more intuitive traction to the case for closure, since we will be concerned with agents who can strengthen the set of reasons on which they are basing the expansion of their belief systems in episodes of reasoning by adding that piece of philosophical information to the set (which may, in turn, reasonably lead to the kind of perplexity made famous by Lewis Carroll's Tortoise).¹³ I don't mean to deny that this does make matters more dramatic. But it is clear that a closure claim involving known implication is irrelevant to the discussion of the cognitive lives of the vast majority of agents. Very few people even have the concept of logical implication! I mean, very, *very* few! It is more than a little ironic that the discussion of closure theses involving knowledge of implication took root in the externalist literature in epistemology. We are used to thinking of externalism as anathema to over-intellectualization in epistemology.

In any case, some of the recent entries in the debate stem from that kind of objection to the older, Dretske-derived claims. Two prime examples:

- (e) "If one knows P and competently deduces Q from P, thereby coming to believe Q, while retaining one's knowledge that P, one comes to know that Q." (Hawthorne 2005; also Dretske 2005, Kvanvig 2005a, Williamson 2000, Vogel 1990);
- (f) "[K]nowing p_1, \dots, p_n , competently deducing q, and thereby coming to believe q is in general a way of coming to know q." (Williamson 2000; also Stine 1976).

Does the argument of the previous section have any impact on closure claims like (e) and (f)?

An important difference in the case of agents who are competent in their deductive practices but do not possess the concept of logical implication is that we may reasonably expect them to believe the Russellian premises. We may also

¹² It is noteworthy that, if we were to adopt Crispin Wright's distinction between closure and transmission (Wright 2003), and if, as I have argued, we must acknowledge the dilemma posed by a paradox like Russell's (again, the dilemma whose first horn is disbelieving the paradoxical premises and whose second horn is embracing dialetheism) as providing us with a reason for withholding belief in epistemic closure claims, we should notice that the paradox may be a genuine counterexample to closure (as characterized by Wright). It cannot be *only* a counterexample to transmission (as characterized by Wright) because it should seem absurd to think that warrant for believing the paradoxical premises may in any way depend on our having a warrant for believing the Russellian contradiction. The case does not obviously conform to his information-dependence template for transmission failure. (As he sees it, every case of closure failure is a case of transmission failure but not conversely. And every case of transmission failure is characterized by him as a case of epistemic circularity — though he makes no use of the term "epistemic circularity".) This is noted in passing. The philosophical worth of Wright's views on closure cannot be gauged here.

¹³ On the Carrollian perils of reflection on one's reasons for belief, the reader is referred to my 2005.

reasonably expect them not to feel compelled to infer (and thus believe) the contradictory conclusion logically supported by those premises. It is easy to imagine that they would simply not move past those premises, thus failing to appreciate the philosophically dire straits in which they would find themselves. Now, would it be reasonable for us to deem their belief in the Russellian premises *unjustified* (or *irrational*)? What would count as good, non-question-begging grounds for thinking them irrational? Surely, not the fact that we, the cognoscenti, are justifiably puzzled! (We certainly need an argument to the effect that one's justifiable puzzlement implies the other's unjustifiable non-puzzlement.) If not that, what then?

Not knowing why one who fails to see the disastrous logical consequences of what he believes on good grounds should be deemed irrational, and failing to see that one could conceivably be justified in believing the Russellian contradiction, I submit that Russell's paradox does seem to refute (that is, to offer a case against) populist closure claims as well.¹⁴

3 Epistemic Closure, that Strange Object of Desire

Heretofore, the anti-closure argument is made for me by three sources. In addition to being crucially helped by Priest's expert testimony and bold claim, I take the smooth ride that's offered me by both the genius of Russell and the genius of Dretske. As we know, it was Dretske who made it possible for all of us to bring epistemic closure claims under severe scrutiny. For epistemologists like us, living under the long shadow of Dretske's 1970 paper, life becomes full of exciting alternatives. In fact, as I hope you will have just seen, where the philosophical logician faces an uphill battle with common sense, we, epistemologists, can bask in the glory of a perfectly satisfactory explanation of the *epistemic* aspect of the phenomenon that occupies the philosophical logician when he ponders the logical paradoxes. (And maybe — just maybe — that's *all* the satisfaction that anybody will *ever* have concerning the toughest logical paradoxes, such as Russell's.) The proposed explanation is all the more satisfactory for being compatible with any persuasive story that may eventually come our way according to which some of Russell's paradoxical assumptions are, indeed, either false or nonsensical. For, if some of those assumptions turn out to be, indeed, either false or nonsensical, we shall only need remind ourselves that we are all *fallibilists*, that truth and rationality are independent (even for those who think rationality truth-conducive), and that a century of philosophical logic has failed to make a dent in our paradoxical intuitions. In the 21st century, life is good for an epistemologist!

¹⁴ I anticipate this objection: "Suppose you take it on good authority that p but fail to notice that p implies that not-p. Surely, it is an objective fact that you shouldn't believe that p." Isn't this ultimately motivated by some form of justification-closure dogmatism? I just can't see that such a far-reaching condemnation of *inferential near-sightedness* is philosophically warranted by a tenable theory of epistemic rationality. I suppose we don't want to count such a *common* form of "stupidity", of inferential mediocrity as a case of *irrationality*. Clearly, there is something undesirable in the cognitive life of such an agent. He can't know that p. But, assuming fallibilism, I can't see a charge of irrationality that will stick. Maybe the objection rests on a confusion between synchronic and diachronic desiderata of epistemic rationality. I will not pursue the matter here. For important admonition in this regard, see Richard Foley (2001).

And good life is ensured by iron-clad attention to how we strayed from the path of thoroughgoing fallibilism back there somewhere. Dretske gave us the opportunity to scrutinize epistemic closure. We've seized the opportunity, and, now, even this very fundamental question, "How plausible is any of the major epistemic closure theses anyway?", is getting some very exciting answers. Consider these claims, and never mind the fact that these authors may not have been thinking about the very same closure thesis:

Dretske 2005: "Yes, closure sounds like an eminently plausible principle."

Harman and Sherman 2004: "Intuitively, it is not generally the case that, if one knows that P and one knows that if P then Q, one should be able to know that Q. [...] Strong closure principles [like this one] are not intuitively plausible [...] and they have no good arguments in their defense."¹⁵

Hawthorne 2005: "While principle [(e) above] is compelling, it is not as *manifestly* obvious as the transitivity of the *taller than* relation...[and] we should work to provide some supplementary considerations that reinforce our commitment to [e]'s intuitive status."

Isn't the discrepancy here *shocking* for those who think that closure is intuitively appealing? And doesn't Hawthorne's claim sound awfully close to the oxymoronic "Closure is intuitive, although it may not *seem* to be"?¹⁶

It's actually to Hawthorne's credit that he *blinked* when nothing short of a resolute "Yes!" would do as an answer, from a closure advocate, in reply to the question of whether there is any intuitively compelling epistemic closure thesis. Such hesitation does encourage closure deniers. It's welcome, *philosophically sound* hesitation from where I look. Like Harman and Sherman, I can't see anything even remotely intuitive about epistemic closure claims. Epistemic closure is a very complicated concept, and one into which we have looked seriously for only about forty years now. But I also believe that there are *weighty philosophical reasons* for thinking that closure principles *should seem intuitive*, and that it behooves the closure denier to identify such reasons and come to terms with them.

We shall shortly see how that may be accomplished (in Sect. 7 below).

¹⁵ As Jonathan Kvanvig (2005) notes, the literature on so-called "closure principles" has stretched the terminology to breaking point. Strictly speaking, some of these claims are not really closure claims, since the property referred to in the antecedent of these conditionals is not the property referred to in the consequent. So, "closure" has really become a metaphor of sorts in the literature. I share in Kvanvig's point about this kind of excessive terminological permissiveness, but I'll be willing to play along and call a "closure claim" anything in the neighborhood of closure.

¹⁶ Granted, the thought that intuitive appeal comes in degrees is not obviously false, but that is cold comfort for Hawthorne. It does seem that any such hesitation makes the talk of intuitive appeal self-refuting. We lose our grip on any useful notion of intuition.

4 A Bogus Distinction with an Impressive Pedigree and Other Irritants

A source of dissatisfaction with the foregoing is the now-popular distinction between single- and multi-premise closure. I have completely ignored it until now. But we can't afford to be cavalier about it.

Let me be frank about this: That is one of the most philosophically perverse distinctions in the recent epistemological literature. It should have remained a technicality when it was brought to the fore in the discussion on closure by Gail Stine (1976). Stine herself saw it as a technicality and advised against giving it any importance. But after Richard Foley (1979) and Peter Klein (1985, 1998) saw fit to keep a distance from multi-premise closure as a way around the lottery paradox, a number of philosophers followed in their footsteps and, frankly, blew the distinction way out of proportion.¹⁷ It's easy to get sidetracked here. But we surely can avoid that hazard. These two quick points should suffice.

First, some influential philosophers who could have jumped on the single-premise bandwagon have remained immune to the appeal of that distinction. Thus, Timothy Williamson (2000), heeding Stine's advice, refused to acknowledge single-premise closure as the case of epistemic closure worth defending. And Richard Feldman (1995) thought it easy to steer clear of the motivation for single-premise closure by simply (though only tacitly) asking his reader to focus on a case of ostensive single-premise closure in which the premise is a conjunction of any number of propositions one cares to include in the implying proposition. The main points in the closure/anti-closure debate, he suggests, seem completely unaffected by considering any conjunction one cares to put together.

Second, when you look at the *intuitive* case for epistemic closure, the closure advocate doesn't seem to find any use for the single-premise/multi-premise distinction — which comes as no surprise, really, since there is absolutely nothing to recommend such a distinction from a purely commonsensical point of view. It is a philosopher's artifice. Consider, for example, this passage in John Hawthorne's case for closure (in which he essentially echoes a passage in Feldman 1995):

Suppose Q is a “heavyweight” consequence of P and S knows P and also that P entails Q.¹⁸ I ask S whether she agrees that P. She asserts that she does [...] I then ask S whether she realizes that Q follows from P. “Yes”, she says. I then ask her whether she agrees that Q. “I'm not agreeing to that”, she says [...]

¹⁷ These include John Hawthorne (2005), Doris Olin (2005), Jason Stanley (2005), Richard Fumerton (2006). Stine rightly suggests that it is philosophically ill-advised to see it as a basis for paradox resolution.

¹⁸ In his 2005 paper, Dretske called “heavyweight implications” those propositions “we needn't know even though we know our knowledge depends on their truth”. He hasn't offered an explicit definition of the term “heavyweight implication”, but the concept is not hard to grasp (not hard for our purposes here anyway). According to some of his examples, one may know that there are cookies in the jar without knowing the heavyweight implication that there is a material world; one may know that there are zebras in the pen without knowing the heavyweight implication that the zebra-looking objects are not cleverly disguised mules. It is typically thought that the class of heavyweight propositions includes the negations of skeptical hypotheses.

interlocutor now resembles perfectly Lewis Carroll's Tortoise, that familiar object of ridicule who was perfectly willing to accept the premises of a *modus ponens* argument but was unwilling to accept the conclusion [...] [Hawthorne (2005), 32]

I have no idea how Hawthorne's outrage could possibly hold sway over anyone who is already persuaded by Dretske's case against closure, since a closure denier may have already granted that closure failure can be perceived as shocking (as Dretske himself seems to do). Some things in life are shocking. But just noticing that they are may not suffice to put them out of existence! In any case, the main feature of the above passage for our purposes here is that an ostensive case for *single-premise* closure (a case for thesis (e)) is put forward under the guise of an aborted *modus ponens* — thus completely missing its ostensive, single-premise mark!¹⁹ It misses its mark for good reason, though: There is no intuitive basis for a single-premise/multi-premise distinction. An intuitive case against the epistemic effectiveness of deduction — as implied by Williamson's view of the matter — is a case against *any* instance of attempted transmission of epistemic status by deduction.

Then there are the ills arising from the current, embryonic state of the epistemology of reasoning *as a recognized field of epistemological inquiry* — specifically, the fact that popular closure claims are usually expressed by conditionals with bloated antecedents — like (e) above — where we find misplaced reference to a number of necessary conditions for the acquisition of knowledge or justification in episodes of reasoning that had better be kept out of the way in a discussion of whether epistemic status is always transmitted by logical implication. Some authors (Klein, Olin, and others) keep those additional points for their comments on the margins about whatever else, in addition to the transmission of epistemic status, may affect the generation of knowledge or justification in episodes of deductive reasoning.²⁰ But most do seem happy with the bloated antecedents. In discussing closure, we should, however, be focusing exclusively on whether epistemic status is always transmitted by logical implication. For illustration, consider this inductivist variant of (e):

(e') If one knows P and competently *inductively* infers Q from P, thereby coming to believe Q, while retaining one's knowledge that P, one comes to know that Q,

and you will notice that we would not want to subtract any of those conditions pertaining to aspects of the inference other than its inductive nature. Which makes the point that all those necessary conditions not having to do with the *logical* aspect of the deductive character of an inference had better be discussed in other chapters of a complete epistemology of reasoning.

¹⁹ For more criticism of Hawthorne's case for closure, see Fred Adams and Julia Figurelly, forthcoming publication.

²⁰ I'm not casual in speaking of the *generation* of knowledge or justification in episodes of reasoning where Robert Audi (2003), for one, would have preferred to speak simply of the *transmission* of epistemic status via reasoning. We must take account of inferential knowledge arising from benign falsehoods. My 2006 offers an introduction to the matter and references for further reading.

With these admonitory points in mind, I will henceforth focus on the following, multi-premise variant of a closure claim favored by Peter Klein (2004), which, to my mind, is the very strongest closure claim we can identify in the literature on the issue.

- (g) If S is justified in believing that P_1, \dots, P_n ($n \geq 1$), and P_1, \dots, P_n jointly imply that Q , then S is justified in believing that Q .

(Still, note that my counterexamples are expected to work against any of the closure claims favored by non-justificationists as well. I just want to take account of a notion of epistemic entitlement that applies to what an agent *may* justifiably believe or know, even if that kind of entitlement is supposed to extend, by implication, to what one may never be nomically able to comprehend given the present state of one's cognitive abilities. The core notion of epistemic entitlement should not be bounded by any contingent limits of human ability.)

But (g) is much too coarse for our purposes. I should like to help myself to a familiar distinction of *prima facie* and *ultima facie* justification, the former being a necessary condition of the latter. The distinction I have in mind is one according to which *prima facie* justification is to be taken as *knowledge-grade justification* — but in a sense of the term that does not rule out its being overridden justification. And *ultima facie* justification simply adds the non-existence of non-overridden overrides in one's mental life to the *prima facie* justification of a given proposition for a given agent.²¹

With that distinction in hand, and reading our main target, (g), as

- (g') If S is *ultima facie* justified in believing that P_1, \dots, P_n ($n \geq 1$), and P_1, \dots, P_n jointly imply that Q , then S is *ultima facie* justified in believing that Q ,

I think some other counterexamples to closure are forthcoming.

5 Another Paradox, Another Refutation

Before looking into those weighty reasons hinted at in Sect. 3 above for a conciliatory move in the debate, let me give you another counterexample to every major epistemic closure claim (but, again, ignoring any attempts to dignify the single-premise/multi-premise distinction, since *only* multi-premise closure matters).

Suppose I justifiably believe or know the following:

- (i) Everything my psychoanalyst says about me is true.
- (ii) My psychoanalyst says that (both) my father loves me but I don't believe my father does.

By all accounts, I'm *not* in the clear to either know or justifiably believe that

- (iii) Therefore, my father loves me but I don't believe he does,

which is, of course, an instance of Moore's paradox.²²

²¹ I adhere to the terminology put forward in Klein (1981) in speaking of "overrides". If you prefer "defeaters", notice that these must be *internal* defeaters.

²² I thank Peter Klein and Roy Sorensen for discussing this case of apparent closure failure with me in 1998.

Although the whole story of why the conclusion of this Moore-paradoxical argument cannot be a case of justified belief *for me* (for the first-person believer in the example) is much too long for the available space, you can readily see (with John Williams) that *my* belief in the conclusion must be false, since, if belief distributes over conjunctions (as it clearly seems to), my belief in the proposition on the left-hand side of the conjunction (“My father loves me”) falsifies the conjunction by way of falsifying the proposition on the right-hand side of it (“I don’t believe he does”).²³ So, plainly, the conclusion cannot be a case of knowledge for me.²⁴ And, henceforth, you can at least justifiably suspect any claim to the effect that I can justifiably believe a proposition that I cannot know, even though that proposition would be deductively justified for me by a set of premises all of which can simultaneously be cases of knowledge for me. So, again, we clearly seem to have found good reason for thinking that we are looking at a valid argument with (by hypothesis) true premises the conclusion of which cannot be the object of justified belief or knowledge.

6 Mundane Closure Failure — Without Heavyweight or Paradoxical Propositions

Consider my third case, *Flight 80*: Suppose I have *ultima facie* justification for believing that

(i) If I take flight 80, I will land in Tel Aviv.

And suppose I also have *ultima facie* justification for believing that

(ii) I have just taken flight 80,

but there is also the same kind of justification for believing that

(iii) There are terrorists on board.

I submit that the non-philosophical rational agent will not feel — and should not feel — at all tempted to perform that *modus ponens* (while believing the premises). But notice that (iii) is perfectly compatible with the fallibilist assumptions according to which the agent may have ordinary, *zebra-grade* justification for belief in (i), knowledge-grade justification of the kind that is familiar to fallibilists of every stripe (no pun intended!). If, on popular fallibilist assumptions, one can know a conditional like (i) at all, that kind of knowledge-grade evidence will not have required any form of nomic necessitation of the consequent by the antecedent. Everybody knows that planes go down every now and then. But, as long as those mishaps are rare enough, we seem to have no problem making knowledge claims like (i). The relevant fallibilist intuitions

²³ Williams’ point is made in his 1994 paper. My explanation of Moore-paradoxical knowledge/justification deprivation takes two forms: the evidentialist one, in my 2001 paper, and the non-evidentialist one, in my 2007 paper on the matter.

²⁴ Commenting on a draft of this passage in correspondence, Williams adds: “Here’s another reason: If you know (iii) then you know its second conjunct, but since whatever you know is true, you lack the belief that your father loves you, which is a necessary condition of knowing that he loves you.” Thank you, John!

have been scrutinized *ad nauseam* in the literature and they do seem to provide a solid foundation for various descendants of the early theories of relevant alternatives.

The natural move against this counterexample, I suppose, will be some attempt to claim that (iii) is *somehow* an overrider of the agent's justification for believing (i). So, let us suppose, to begin with, that (iii) is what John Pollock and Joseph Cruz (1999) have called "a rebutting defeater" — roughly, a reason (however weak it might be) for believing that (i) is false.

My reply to that relies on two claims that are hardly objectionable:

Evidential distribution (ED): If E is good evidence for believing that P&Q (for a given agent at a given time), then E is just as good evidence for believing that P (for that agent at that time) and just as good evidence for believing that Q (same agent and time).²⁵

and

Evidential contradictoriness (EC): If E is some (not necessarily good) evidence for S to disbelieve that P (to *reject* P) at t, then E is just as good evidence for S to believe that not-P at t.²⁶

If you grant me these two theses, it becomes easy to see that, if (iii) is an overrider (of the rebutting kind) of one's justification for (i), then, by (EC), it must be some reason to believe the negation of (i). But the negation of (i) is logically equivalent to (i') I have taken flight 80, but I will not land in Tel Aviv.

And we may safely assume that (i') is exactly what the ordinary believer would take the contradictory of (i) to be in the envisaged circumstances. (There is no risky assumption about the meaning of (i) here.) However, by (ED), (iii) would have to be some reason for believing that I have taken flight 80, which is absurd!

Next, one may try to characterize (iii) as what Pollock and Cruz have called "an undercutting defeater" — roughly, a reason for thinking that one's reasons for believing (i) are not good enough. Thus, one may hold that the link immediately preceding (i) in my evidential chain leading to a justified belief in (i) includes the belief that (iv) *all is well with flight 80*. That is, it may be held that my good inductive evidence for pleasant thoughts about flight 80 supports (iv), which, in turn, plays an essential role in providing me with a justification for believing (i). It may, then, be held that (iii) implies (together with whatever unobjectionable assumptions we might have to add for logical implication) that (v) *something is not well with flight 80*. Now, it should seem that I can't justifiably believe all of (i)–(iii), since a logical consequence of (iii) would override my justification for believing (i) (by contradicting (iv)).²⁷

²⁵ If I'm not mistaken, this thesis was first introduced in the literature in Klein 1981 (though I phrase it differently and give it a label here).

²⁶ Although this is certainly very close to the surface in the literature and tacitly assumed by many, I can't recall any statement of it. (EC) is a partial explanation of the concept of *counterevidence*.

²⁷ I thank Doris Olin for calling my attention to the need to discuss this objection, though she may not approve of the way I have fleshed out her sketch of the objection.

The objection is undermined by two points. First, in assuming that one's evidence for (i) *must* include the false (iv), we do seem to be letting the infallibilistic demand for nomic implication of the consequent of (i) by its antecedent creep into the picture. On the most natural construal of the objection, my knowing (i) depends on my knowing (iv). Why is that a necessary condition on knowledge of (i)? The only answer in sight seems to be a requirement of nomic necessitation for conditional knowledge: you don't know a conditional unless the truth of the antecedent nomically guarantees the truth of the consequent. One of the problems with this kind of requirement is that it is incompatible with the view that one may have knowledge on the basis of purely statistical evidential support. This is a problem because the kinds of beliefs giving rise to both the lottery and the preface paradoxes can be seen as instances of rational believing on the basis of purely statistical evidence (even though, unlike the lottery, the preface paradox does not essentially involve statistical evidence).²⁸ (And, if it is held that statistical evidence is sufficient for rational belief but not for knowledge, it will certainly take a very complicated explanation for us to understand how that can be.) In the case of flight 80, if statistical support is not ruled out as a basis for rational belief and knowledge, it becomes reasonable to hold that one retains knowledge of a conditional like (i) while losing one's rational basis for believing the consequent of (i) (in view of (iii)), since, in setting up the case, we can make the favorable ratio of successful to unsuccessful landings for flight 80 as statistically impressive as we want! Granted, this is not a decisive rebuttal of the objection, but it does show that the view on which the objection is based may be prohibitively costly.²⁹

Second, and decisively, it is absurd to think that (iv) *must be* (in any shape or form) in the evidential ancestry of (i). Presumably, I may come to know (i) from believing it on the basis of reliable testimony — which would clearly seem to obviate the need for any additional assumptions about flight 80 on my part.

It doesn't look like a good case will emerge for the view that (iii) is an overrider of my justification for believing (i). So, why not closure failure?

Lastly, consider *The Conference*: Suppose you have *ultima facie* justification for believing that

(i) Whenever C is invited to speak at the Bled Conference, C is ecstatic.

And suppose you also have *ultima facie* justification for believing that

(ii) On August 14, 2006, C is invited to speak at the Bled Conference.

Now, suppose that you have that same kind of non-overridden justification for believing that

(iii) On August 13, 2006, C's mother died (and C loved his mother dearly, and C was immediately aware of her death, and C has excellent memory, etc.).

²⁸ For criticism of the view that rejects purely statistical evidence as a basis for rational belief, see David Christensen (2004), Chap. 3, and Doris Olin (2003), Chap. 5. On the relevant difference between the lottery and the preface, see Earl Conee (1992).

²⁹ It's certainly uncontroversial that every reasonable fallibilist must feel some measure of attraction to David Lewis' (1996) claim that fallibilism is a form of madness. Discussing that fundamental tension is outside the scope of this paper. We're merely exploring a fallibilist alternative to the present difficulties.

You don't think you would be *ultima facie* justified in believing that C was ecstatic on August 14, do you? And, again, I see no basis for denying that we can have zebra-grade knowledge of (i). Nor do I see how (iii) can plausibly be regarded as an overrider of the believer's justification for believing (i).

Again, one may — perhaps reasonably — shy away, on skeptical grounds, from believing *any* empirical generalization. We've long known that that theoretical option is open to us. But, as Dretske has noted, we want to explore our non-skeptical options. A main point of this paper is that closure denial is motivated without appeal to skeptical ("heavyweight") considerations.

7 Not Every Closure Failure Succeeds

Now, notice that none of the foregoing is any reason to think that the following thesis is false:

- (h) If S is *ultima facie* justified (to degree n) in believing that P_1, \dots, P_n ($n \geq 1$), and P_1, \dots, P_n jointly imply that Q, then S is *prima facie* justified (to degree n) in believing that Q.³⁰

Every putative counterexample I have offered above — including Russell's paradox and the one derived from Moore's paradox — is a case in which there seems to be a non-overridden overrider of one's justification for belief in the implied proposition.³¹ None of my counterexamples carries any weight against (h). This is an important feature of my case against closure. It gives me the opportunity to part company with Dretske as regards one of his boldest — and, to my mind, ultimately indefensible — claims.

It will be recalled that one of the most distinctive features of Klein's case for closure is his claim that the implied proposition may get *all* of its evidential support from the implying proposition. Dretske (2005) has claimed that this is "verbal hocus pocus": "We have our cake and eat it too", Dretske writes about the Kleinian view of the matter, "by simply *stipulating* that we know whatever we know to be implied by what we know even though we have no identifiable way of knowing it" (Dretske 2005, emphasis added). Though this is, in essence, what he has held for over 35 years, Dretske may have derived some encouragement from Anthony Brueckner, who has recently claimed that, "if the entailing justified proposition that the animal

³⁰ Robert Audi (1991) suggests that a weaker principle might be defensible: "[I]f S has reason to believe p, and S can understand both q and the proposition that it is entailed by p, then S has reason to believe q". This is most significantly different from (h) for not concerning *knowledge-grade* evidence.

³¹ In the case of the paradoxical counterexamples, the non-overridden overrider may be as simple as this: "But that's absurd!" The overrides in the other cases are given in the text. It should also be noted that, for the Dretskean counterexamples (the cases involving "heavyweight" implications), the non-overridden overrides are the conjunctions encompassing every claim in a skeptical argument — or maybe just the conclusion of the argument — of which the believer is aware but is rationally unable to deny.

is a zebra is *not* itself a further source of justification (as I have argued against Klein), then the entailed proposition has no adequate source of justification at all” (Brueckner 2000).

Now, I can’t review Brueckner’s argument here, but allow me to suggest that the above Dretskean view (the “verbal hocus pocus” view) *must be wrong*. It leaves us no room at all to understand the alleged intuitive appeal of closure claims. But any view which goes that far has certainly gone way too far to be fully acceptable. Richard Feldman’s (1995, 1999) oft-quoted expressions of outrage about claims of closure failure have clearly struck a nerve in the literature, and I am one of those who think that *something* is clearly right about that kind of outrage, though it’s hard to see how *anything* can be right about that on Dretske’s account of the matter.

Klein’s distinctive claim in the dispute over closure simply acknowledges the time-honored view that entailing reasons are, in a sense, *conclusive reasons*: knowledge-grade justification in, knowledge-grade-justification out. Thesis (h) captures just that. If logical implication did not transmit knowledge-grade justification, it would be an utter mystery why implying evidence has, since time immemorial, been recognized as the best evidence anybody can have for anything. We would have no explanation for “logical compulsion”. This is what closure supporters are justifiably trying to capture.

8 Concluding Remarks: Defeasibility and the Myth of Conclusive Reasons

But notice how principle (h) is anathema to any important, historical notion of the *conclusiveness* of reasons: It obviously is *oxymoronic* to speak of having *conclusive reasons* whose evidential effectiveness may be (ultimately) overridden! Consider the following claims about how implying reasons are *conclusive* reasons — reasons the evidential power of which simply cannot be destroyed. According to John Pollock and Joseph Cruz (1999, 36–7),

[w]e shall call a reason that logically entails its conclusion a *conclusive reason*. Inductive reasons are nonconclusive reasons [...] The most important characteristic of nonconclusive reasons is that they are *defeasible* [...] Defeasible reasons are reasons for which there can be defeaters. Such a reason is called a *defeasible reason*.

It is very clearly (contextually) implied here that entailing reasons are not defeasible.

Here’s a similar passage in Doris Olin (2003, 9):

Valid arguments [...] do not share with inductive arguments the feature of defeasibility. Given premises we are entitled to accept as true, valid reasoning will yield a conclusion we are entitled to accept as true, *no matter what further information we have*.

Here is the same old doctrine of the conclusiveness of entailing reasons as endorsed by Matthias Steup (1996, 12):

[B]ecause deduction is truth-preserving, it allows us to start out with premises for which we have conclusive evidence and to end up with conclusions for

which we have conclusive evidence. Conclusive evidence for a conclusion cannot be achieved with nondeductive types of derivation.

Or, again, consider this passage in Michael Williams (2001, 45), where monotonicity is apparently mixed up with epistemic conclusiveness (that is, the non-existence of overrides, here called “defeaters”):

Because validity cannot be compromised by adding premises, deductive arguments are sometimes said to be *monotonic* [...] If a given body of evidence logically entails a particular conclusion, it is absolutely conclusive in that it excludes every possible ‘defeater’ for that conclusion: that is, every possible situation in which that conclusion would be false.

These authors are all saying essentially the same thing.³² And what they are saying is *false*: There are no *conclusive reasons* in this traditional sense of the term. Every form of reasoning is *epistemically defeasible*. The failure of epistemic closure was easy to overlook until Dretske showed us the way out of a confusion of millennial proportions between monotonicity — closure of truth under logical implication — and epistemic closure. We can finally let go of this last remnant of infallibilism in epistemology.

³² Count Mark Sainsbury (2001, pp. 11–12) among those who have explicitly subscribed to the traditional view of implying reasons as conclusive reasons.

Acknowledgments I have benefited from insightful comments on earlier drafts of the paper generously offered by Fred Adams, Jonathan E. Adler, Rodrigo Borges, Anthony Brueckner, Stewart Cohen, Rogel de Oliveira, Fred Dretske, Patrick Greenough, Nicholas Griffin, Peter Klein, Jonathan Kvanvig, Nenad Miscevic, Ernest Sosa, Marcus Willaschek, and John N. Williams. I am most especially indebted to Doris Olin and Danilo Suster for much helpful discussion. A draft of the paper was presented at the 2007 Bled Conference on Contemporary Epistemology. I am grateful to my audience there.

References

- Adams, F., & Figuerely, J. Towards closure on closure. (in press).
- Audi, R. (1991). Justification, deductive closure, and reasons to believe. *Dialogue*, 30, 79–82.
- Audi, R. (2003). *Epistemology: A contemporary introduction to the theory of knowledge* (2nd ed.). New York: Routledge.
- Brueckner, A. (2000). Klein on closure and skepticism. *Philosophical Studies*, 98, 139–151.
- Canary, C., & Odegard, D. (1989). Deductive justification. *Dialogue*, 28, 305–320.
- Christensen, D. (2004). *Putting logic in its place: Formal constraints on rational belief*. Oxford: Oxford University Press.
- Cohen, S. (1988). How to be a fallibilist. In J. Tomberlin (Ed.), *Philosophical Perspectives 2* (pp. 581–605).
- Conee, E. (1992). Preface paradox. In J. Dancy & E. Sosa (Eds.), *A companion to epistemology*. Oxford: Blackwell.
- De Almeida, C. (2001). What Moore's paradox is about. *Philosophy and Phenomenological Research*, 62, 33–58.
- De Almeida, C. (2005). The paradox of inference in canonical form, *Certain Doubts*, Baylor University, J. Kvanvig (ed./mod.), http://fleetwood.baylor.edu/certain_doubts/?p=499.
- De Almeida, C. (2006). Testimony, lies, fiction and benign falsehoods, *Certain Doubts*, Baylor University, J. Kvanvig (ed./mod.), http://fleetwood.baylor.edu/certain_doubts/?p=639.
- De Almeida, C. (2007). Moorean absurdity: An epistemological analysis. In M. Green & J. N. Williams (Eds.), *Moore's paradox: New essays on belief, rationality, and the first person* (pp. 53–75). Oxford: Oxford University Press.
- Dretske, F. (1970). Epistemic operators. *Journal of Philosophy*, 67, 1007–1023.
- Dretske, F. (2005). The case against closure. In M. Steup & E. Sosa (Eds.), *Contemporary debates in epistemology* (pp. 13–26). Oxford: Blackwell.
- Feldman, R. (1995). In defense of closure. *Philosophical Quarterly*, 45, 487–494.
- Feldman, R. (1999). Contextualism and skepticism. In J. Tomberlin (Ed.), *Philosophical Perspectives 13* (pp. 91–114).
- Foley, R. (1979). Justified inconsistent beliefs. *American Philosophical Quarterly*, 16, 247–257.
- Foley, R. (2001). *Intellectual trust in oneself and others*. Cambridge: Cambridge University Press.
- Fumerton, R. (2006). *Epistemology*. Oxford: Blackwell Publishing.
- Harman, G. & Sherman, B. (2004). Knowledge, assumptions, lotteries. In E. Sosa & E. Villanueva (Eds.), *Philosophical Issues 14* (pp. 492–500).
- Hawthorne, J. (2005). The case for closure. In M. Steup & E. Sosa (Eds.), *Contemporary debates in epistemology* (pp. 26–43). Oxford: Blackwell Publishing.
- Klein, P. (1981). *Certainty: A refutation of scepticism*. Minneapolis: University of Minnesota Press.
- Klein, P. (1985). The virtues of inconsistency. *Monist*, 68, 105–135.
- Klein, P. (1998). Knowledge, concept of. In E. Craig (Ed.), *The Routledge encyclopedia of philosophy*. London: Routledge.
- Klein, P. (2004). Closure matters: Academic skepticism and easy knowledge. In E. Sosa & E. Villanueva (Eds.), *Philosophical Issues 14* (pp. 165–184).
- Kvanvig, J. (2005). The closure mess, *Certain Doubts*, Baylor University, J. Kvanvig (ed./mod.), http://fleetwood.baylor.edu/certain_doubts/?p=347.
- Kvanvig, J. (2005a). Denials of closure?, *Certain Doubts*, Baylor University, J. Kvanvig (ed./mod.), http://fleetwood.baylor.edu/certain_doubts/?p=311.
- Lewis, D. (1996). Elusive knowledge. *Australasian Journal of Philosophy*, 47.4, 549–567.
- Olin, D. (2003). *Paradox*. Montreal: McGill-Queen's University Press.
- Olin, D. (2005). A case against closure. In C. de Almeida (ed.), *Perspectives in contemporary epistemology*, Veritas, 50/4 (Porto Alegre, Brazil: EDIPUCRS), <http://revistaseletronicas.pucrs.br/veritas/ojs/index.php/veritas/issue/view/178>.

- Pollock, J., & Cruz, J. (1999). *Contemporary theories of knowledge*. Lanham: Rowman & Littlefield.
- Priest, G. (1986). Contradiction, belief and rationality. *Proceedings of the Aristotelian Society*, 86, 99–116.
- Priest, G. (2006). *In contradiction: A study of the transconsistent* (2nd ed). Oxford: Oxford University Press.
- Pritchard, D. (2005). *Epistemic luck*. Oxford: Oxford University Press.
- Roush, S. (2005). *Tracking truth*. Oxford: Oxford University Press.
- Sainsbury, M. (2001). *Logical forms* (2nd ed). Oxford: Blackwell.
- Stanley, J. (2005). *Knowledge and practical interests*. Oxford: Oxford University Press.
- Steup, M. (1996). *An introduction to contemporary epistemology*. Upper Saddle River: Prentice Hall.
- Stine, G. (1976). Skepticism, relevant alternatives, and deductive closure. *Philosophical Studies*, 29, 249–261.
- Vogel, J. (1990). Are there counter-examples to the closure principle?. In M. D. Roth & G. Ross (Eds.), *Doubting* (pp. 13–27). Dordrecht: Kluwer.
- Williams, J. (1994). Moorean absurdity and the intentional ‘structure’ of assertion. *Analysis*, 54, 160–166.
- Williams, M. (2001). *Problems of knowledge*. Oxford: Oxford University Press.
- Williamson, T. (2000). *Knowledge and its limits*. Oxford: Oxford University Press.
- Wright, C. (2002). (Anti-)sceptics simple and subtle: In G. E. Moore & J. McDowell (Eds.), *Philosophy and Phenomenological Research*, 65/2.
- Wright, C. (2003). Some reflections on the acquisition of warrant by inference. In S. Nuccetelli (Ed.), *Semantic externalism and self-knowledge*. Cambridge: MIT Press.