

Wright Back to Dretske, or Why You Might as Well Deny Knowledge Closure

MARC ALSPECTOR-KELLY

Western Michigan University

Fred Dretske notoriously claimed that knowledge closure sometimes fails. Crispin Wright agrees that warrant does not transmit in the relevant cases, but only because the agent must already be warranted in believing the conclusion in order to acquire her warrant for the premise. So the agent ends up being warranted in believing, and so knowing, the conclusion in those cases too: closure is preserved. Wright's argument requires that the conclusion's having to be warranted beforehand explains transmission failure. I argue that it doesn't, and that the correct explanation does not imply that the agent will end up warranted in believing the conclusion when transmission fails. Those who agree that transmission does fail in those cases, therefore, might as well follow Dretske in denying knowledge closure too.

“Here’s the way to think about closure. Sometimes we know the consequent by inferring it from the antecedent and sometimes we know it prior to knowing the antecedent. But however we know the consequent, it remains impossible to know the antecedent without being in a position to know the consequent because if one did not already know the consequent, one could still infer it.”
— Cohen (2000, 138).

“[O]nce one appreciates the wholesale failure of evidential transmission, the failure of closure is, if not mandatory, easier to swallow.” — Dretske (2005, 15).

I. Formulation

Knowledge closure is, or at least is in the vicinity of, the following claim:

KC: Necessarily, for every agent S and propositions P and Q, if (i) S knows P, (ii) competently infers Q from P while knowing P throughout, and (iii) believes Q because she has just so inferred, then S knows Q.

With a slight modification, this is John Hawthorne's formulation of closure, itself inspired by Timothy Williamson.¹ Hawthorne's version has, to a significant extent, superseded earlier formulations in which (ii) was preceded by "knows that P implies Q" and (iii) by "believes Q", which formulations faced various problems avoided by Hawthorne's formulation.² I will explain the slight modification before proceeding.

In Hawthorne (2004) clause (iii) reads "thereby comes to believe Q". This excludes cases in which the agent already believes Q but does not yet know it but who, as a consequence of having performed the inference from P, comes to know what she previously only believed.³ These are surely cases that should be covered: closure would intuitively be violated if such an agent only ended up believing Q but did not know it. Clause (iii) in KC avoids this problem so long as beliefs can be causally overdetermined. The agent can then believe Q because she inferred it from P compatibly with her already believing Q on other grounds.

II. Closure and Transmission

Few epistemologists deny closure. Its advocates typically treat it as an epistemological axiom, one so obviously correct that accounts of knowledge that imply its falsehood—famously the classical sensitivity accounts originally formulated by Fred Dretske and Robert Nozick—are seriously undermined if not refuted outright as a result.⁴

Behind the seeming obviousness of closure lies a persuasive intuition concerning the power of competent deductive inference to generate knowledge from known premises. The Hawthorne/Williamson approach to closure's formulation is motivated precisely by the intuition that competent deductive inference is an infallible way to extend one's knowledge.⁵

¹ Hawthorne 2004, 34; Williamson 2000, 117. KC is, however, more distant from Hawthorne's later formulation in his 2005. In fn. 7 below, I will explain the reason for preferring his earlier version. KC concerns in particular single-premise closure, upon which this paper focuses for simplicity's sake.

² For example, an agent could know P and that P implies Q, and believe Q, but not believe Q because it follows from P, believing Q instead as a result of wishful thinking.

³ Hawthorne recognizes that his formulation is silent on such cases, notwithstanding the plausibility that they are naturally covered by the closure principle. See Hawthorne 2004, 34, fn. 86. Also see fn. 7 below.

⁴ "Robert Nozick's counterfactual analysis of knowledge is famously inconsistent with intuitive closure, but that is usually taken as a reason for rejecting the analysis, not for rejecting closure". Williamson 2000, 117. Dretske started denying closure in Dretske 1970; Nozick did so in his 1981.

⁵ "Williamson has an insightful take on the root of epistemic closure intuitions, namely the idea that 'deduction is a way of extending one's knowledge'." Hawthorne 2005, 41, fn. 6, quoting Williamson 2000, 117.

This intuition is, however, really about transmission. Knowledge transmission takes place when knowledge of the conclusion is acquired as a result of competent performance of an inference from known premises.⁶ KC does not require this; it requires only that the agent somehow ends up knowing the conclusion.⁷

Performance of the inference is obviously not responsible for the truth of the conclusion when knowledge transmits. Since the inference is valid and the premise true, the conclusion is true whether or not the inference is performed. Nor does such a performance secure knowledge because it secures belief. Clause (iii) of KC is included precisely because an agent could perform the inference—or at least recognize its validity—and nevertheless fail to believe the conclusion.⁸

The intuition—and clause (ii) of KC that expresses it—concerns, rather, the transmission of warrant in Alvin Plantinga’s sense, being that which makes for the difference between mere true belief and knowledge.⁹ This implies no theory of warrant. Warrant may or may not imply belief or truth, and it may be external or internal or have elements of both.¹⁰ Plantinga-warrant—“P-warrant”—is merely a placeholder for whatever it is that takes a truly believed proposition to one that is known.

Warrant for Q, unlike knowledge of Q, permits multiple realization. I can’t acquire new knowledge of Q when I know Q already.¹¹ But I can acquire a new warrant for Q when I possess a warrant for Q already. Warrants are essentially ways of knowing, and I might well know Q in more than one way.

Antecedent knowledge of, and so possession of a warrant for, the conclusion does not, therefore, stand in the way of the acquisition of an

⁶ As this characterization indicates, all cases that satisfy transmission also satisfy closure.

⁷ Instead of his earlier “S knows Q”, the consequent in Hawthorne’s later 2005 runs “one comes to know that Q”. If what Hawthorne intends is equivalent to “one comes to know thereby that Q” (and it is hard to see what it would mean otherwise), then the later formulation conflates closure and transmission. But as indicated here, these are distinct. It is moreover important for the closure advocate to distinguish them since there may be cases in which transmission fails without thereby threatening closure. Oddly, Hawthorne is well aware of this, and points out in Hawthorne 2004, 34, fn. 86 that his version of closure there allows for cases in which transmission fails. While he is unconvinced by purported cases of transmission failure, he recognizes that the possibility that there are cases in which transmission fails but closure succeeds should not be foreclosed by the formulation of closure itself. And yet that is precisely what his later formulation in Hawthorne 2005 appears to do.

⁸ “S infers Q from P” might be thought to be conceptually equivalent to “S acquires belief Q as a result of S’s recognition that it follows from P”, so that one can’t infer to Q without believing it. If so, the “recognition of validity” formulation should do.

⁹ Plantinga 1993a and 1993b.

¹⁰ See, however, footnote 40 below.

¹¹ Unless, of course, I somehow lost that knowledge in the interim.

additional warrant by inference. If competent deductive inference really has the power ascribed to it by the intuition, then an agent who performs such an inference from a known premise will thereby acquire a new, previously unpossessed warrant for the conclusion, whether or not she knew it already.

The intuition can, then, be encoded thus:

Necessarily, for every agent *S* and propositions *P* and *Q*, if (i) *S* is warranted in believing *P* and (ii) competently infers *Q* from *P* while being warranted in believing *P* throughout, then *S* acquires a previously unpossessed warrant for *Q* in virtue of having done so.¹²

Call this the “warrant transmission postulate” (WTP). WTP implies that an agent who knows *P*, performs a competent inference to *Q* (while knowing *P* throughout) and believes *Q* as a result, knows *Q*. If she didn’t already know *Q* (because she did not have a warrant for it) then, as Cohen suggests in the first of this essay’s opening quotations, the inference will deliver one to her. So WTP implies KC.

III. Wright and Transmission Failure

Crispin Wright has long argued that warrant transmission sometimes fails (and so that WTP is false): sometimes an agent cannot acquire a warrant by competent inference from a warranted premise. The cases he has in mind include standard putative illustrations of closure failure. In Dretske’s zebra case, for example, Bob, who is visiting a zoo, looks into the zebra paddock, sees an appropriately striped animal, acquires the belief that it’s a zebra, and attempts to infer that it’s not a mule cleverly disguised to look like a zebra.¹³

¹² Transmission is sometimes characterized as requiring that the warrant for the premise is itself carried through the inference to the conclusion, so that the very same warrant for the premise becomes a warrant for the conclusion. Dretske, for example, so characterizes transmission in Dretske 2005, 15. But this is stronger than transmission—at least, it is stronger than I intend the term here. Dretske argues that a reason for believing *P* does not necessarily transmit through the inference so as to constitute a reason for believing *Q*: I might see that there is wine in the bottle without seeing that there isn’t colored water in the bottle, although the former implies the latter. This and similar “penetration” failures provide, he claims, reason to deny that knowledge inevitably penetrates through inference and so that closure fails. But as Klein 1995, Luper 2006 and others point out, the result of inferring from *P* to *Q* is an inference-based belief in *Q*: the reason for believing *Q* is that it follows from *P*, rather than the perceptual reason grounding belief in *P* itself. So failure of transmission in Dretske’s sense is therefore unsurprising, and indicates nothing about whether knowledge will transmit in the sense I employ here (and therefore nothing about closure overall). I’m sympathetic; the argument of this paper is unrelated to Dretske’s penetration-failure argument. Transmission as I understand it requires only that a warrant—but not, or not necessarily, the same warrant—is generated by the inference from *P* to *Q*.

¹³ Dretske 1970.

Dretske famously claimed that, while Bob does know that it's a zebra, he does not know that it's not a cleverly disguised mule, and so knowledge closure fails. Wright interprets such cases instead as instances of warrant transmission failure: Bob cannot acquire a warrant for that proposition by inference from his warranted belief that it's a zebra.

Warrant transmission failures, if they exist, are not necessarily knowledge closure failures. It could be that, while I cannot legitimately acquire a warrant for Q from P by inference, I will end up knowing Q anyway. It could even be that the nature of transmission failure itself is such as to ensure this: the manner in which transmission fails when it does, perhaps, is such as to ensure that I will end up knowing Q.

This is Wright's view.¹⁴ In the cases with which he is concerned, he claims that the agent must be initially warranted in believing the conclusion in order for her to acquire warrant for the premise in the manner in which she acquired it.¹⁵ He claims further that, because antecedent warrant for the conclusion is so required, no new warrant can be transmitted by the inference itself. Transmission therefore fails.

In the zebra case, for example, Bob's acquiring the warrant he has for his belief that it's a zebra—by noting the animal's zebra-esque appearance—is only possible against the background of his already being warranted in believing that it's not a disguised mule. Under the circumstances, if he were not warranted in believing the latter, he would not be warranted in believing the former. As a result, Wright claims, warrant transmission fails.

So while Dretske was right that Bob cannot learn that it's not a disguised mule by inference from it's being a zebra, he was wrong to characterize this as a case of knowledge closure failure. Indeed, the very reason why transmission fails—prior warrant for the conclusion is a precondition of warrant for the premise—ensures that knowledge closure succeeds.¹⁶

¹⁴ Modulo the interpretation of “warrant”, and Wright's recent shift in focus from transmission of warrant to transmission of claims to warrant. See sections IV and V below.

¹⁵ Note that this is not to say that the agent must have a warranted belief in the conclusion. One might possess a warrant for a belief that is nevertheless not based upon that warrant (or any other) and so which belief is itself unwarranted. In such a case the agent has a propositional warrant, but not a doxastic warrant, for her belief. The present issue concerns propositional warrants (although ultimately such warrants must presumably ground belief, and so become doxastic warrants, if those beliefs constitute knowledge).

¹⁶ Silins 2005, fn. 3 lists textual evidence to the effect that this is Wright's view, describing it as the “basic argument” for transmission failure in zebra-style cases. See also the passage from Wright 2011 quoted in section XV below. However, some of Wright's more recent commentary makes this attribution somewhat more tendentious; see section XX.

IV. P-Warrant and W-Warrant

Unfortunately, Wright's sense of the term "warrant", which underlies much of the literature addressing transmission, is not obviously equivalent to P-warrant. While Plantinga uses the term to refer to whatever it takes to convert true belief into knowledge, Wright seems to have in mind something more akin to justification. His most detailed characterization of warrant of which I am aware occurs at Wright (2007, 6):

Let a warrant for a belief be, roughly, an all-things-considered mandate for it: to possess a warrant for P is to be in a state wherein it is, all things considered, epistemically appropriate to believe P. Here I shall require no more specific an understanding of the notion of warrant than that. So the reader is free to construe it further as she thinks fit, whether *externally* (so that, for example, warrant may be constituted by a belief's formation by means of a *de facto* normally reliable cognitive mechanism) or *internally* (so that, for example, warrant may be constituted by a state of information which may be ascertained by a priori reflection and self-knowledge alone), or in more complicated (perhaps admixed) ways.

Call this "W-warrant", to distinguish it from P-warrant.

It is possible for W-warrant, so characterized, to fail to take one from true belief to knowledge in at least three ways. First, the W-warrant in virtue of which it is epistemically appropriate for one to believe P might nevertheless fail to satisfy a higher standard associated with knowledge. Advocates of a probabilistic threshold conception of justification, for example, might claim that the probability of P given one's evidence must be higher in order for one to count as knowing P than is required for one's belief to be justified. So if W-warrant is justification, then one might be W-warranted in believing P without satisfying the more demanding standard required for P-warrant (which delivers knowledge when the belief is true).

Second, one might be Gettiered. Assuming that victims of Gettier scenarios do indeed have unknown but justified true beliefs, they are not P-warranted. So unless a belief can be epistemically inappropriate in Wright's sense merely in virtue of having been Gettiered—which is not the sense one gets from the quotation above or from Wright's work on the issue in general—then victims of Gettier scenarios are W-warranted but not P-warranted.

Third, one could interpret W-warrant as justification while denying that justification is a constituent of knowledge. Perhaps, for example, preference for justified beliefs is in place, not because justification is part of knowledge, but because satisfaction of that preference tends to bring about P-warranted beliefs (at least in conducive circumstances), where P-warrant requires something else entirely.

Suppose, for example, P-warrant is sensitivity: if the belief were false, the agent would not believe it. A belief's manifesting that modal property

implies that the belief is true in fact.¹⁷ But justification, however characterized, is widely regarded as fallible in at least the sense that it is conceptually possible for a belief to be justified and yet false. So sensitivity (P-warrant) and justification (W-warrant) are distinct. Perhaps a belief's being justified increases the likelihood that the belief is sensitive without being part of, or implied by, the sensitivity of the belief thereby produced. A belief could therefore at least be W-warranted without being P-warranted (justified but insensitive) and, perhaps, vice versa (sensitive but unjustified).

Even if W-warrant is closed, therefore, that does not imply that P-warrant is closed as well. In the zebra case, for example, perhaps Bob needs some background evidence in favor of the proposition that the animal in the zebra paddock is unlikely to be a disguised mule in order for him to know that it's a zebra on the basis of its appearance. That evidence may nevertheless be inadequate for P-warrant: while it might render the belief that it's not a disguised mule epistemically appropriate, it might not deliver the where-withal required by knowledge.

Indeed, Dretske pointed this out when introducing the zebra case:

Granted, the hypothesis [that it's a disguised mule] (if we may call it that) is not very plausible, given what we know about people and zoos. But the question here is not whether this alternative is plausible, not whether it is more or less plausible than that there are real zebras in the pen, but whether you know that this alternative hypothesis is false.¹⁸

Whatever one might think background knowledge about people and zoos provides in the way of warranted, justified, appropriate, or reasonable belief, the intuition that you do not possess enough evidence to count as knowing that this animal in particular is not a disguised mule—not having examined the animal's coat at close proximity, and lacking the discriminatory abilities of a zoologist—is very strong.

So the relevance of the literature on transmission underwritten by Wright's sense of "warrant" to the issue of knowledge closure—which concerns P-warrant rather than W-warrant—is uncertain. Nevertheless, in order to ensure its relevance, I will assume that W-warrant is (or implies) P-warrant, so that Wright's views concerning transmission of W-warrant transfer as views concerning transmission of P-warrant, and therefore of knowledge, whether or not this is in fact his intention.

So understood, Wright's view is that P-warrant transmission fails, at least in the cases he considers, because the agent must be P-warranted in

¹⁷ Otherwise, the actual world would be the nearest world in which the belief is false. Since the agent actually believes P, the belief would then be insensitive.

¹⁸ Dretske 1970, 10–16.

believing the conclusion in order to acquire a P-warrant for, and so know, the premise in the manner in which that warrant was in fact acquired.

If that's so then KC follows. Its failure would require that clauses (i)–(iii) of its antecedent be true and its consequent false. If Wright is correct, satisfaction of (i)—that S knows P—requires a P-warrant for Q in these cases. It also requires that Q be true (since P, which is known and so true, implies it). Admittedly, possession of P-warrant for Q does not obviously imply that the agent believes Q. So having a P-warrant for Q might not imply that S knows Q. However, clause (iii)—requiring that S believe Q because she inferred it from P—does imply that S believes Q. So satisfaction of (i) and (iii) together requires that S have a true, P-warranted belief in Q, and so that S knows Q.¹⁹ The consequent will therefore be true as well, and knowledge closure is preserved.²⁰ Call this the Wrightean strategy in defense of knowledge closure.

I will drop the “W-” and “P-” prefixes and simply use “warrant” to reference P-warrant below, except when the difference is relevant to the discussion, at which points I will reintroduce the prefixes.

V. Warrant and Claims to Warrant

Another interpretive concern arises from Wright's more recent work on the issue. In especially Wright (2007), (2011) and (2012), he presents the issue with which he is dealing, not as the transmission of warrant per se, but the transmission of legitimate claims to warrant. He recognizes that, once such a distinction is in place, the relevance of the latter for the former is an open

¹⁹ There is, in fact, another lacuna. Although S is P-warranted in believing Q, and believes it, the latter belief might not be grounded in the former warrant (the inference is not, after all, the source of S's warrant for Q). But knowledge plausibly requires doxastic P-warrant, and therefore such grounding. There are three options. First, insist—implausibly—that knowledge only requires belief and propositional warrant. Second, strengthen Wright's position. While S's being P-warranted in believing P requires only a propositional warrant for Q, knowledge of P requires a doxastic warrant for Q. Third, weaken the consequent of KC to “S is P-warranted in believing Q”. Closure is indeed sometimes formulated in terms of the agent's “being in a position to know” Q; see the quotation from Cohen introducing this essay, for example. Perhaps being in such a position amounts to possession of a propositional P-warrant. Wright's comments at 2012, 458 suggest that he would favor this latter revision. The choice won't matter for what follows, so I will ignore this complication.

²⁰ Perhaps clause (ii)—requiring that S perform a competent inference from P while knowing P throughout—cannot be satisfied in such cases: since warrant transmission fails, one might argue, such an inference can only be incompetent. If this is correct, then Wright's characterization of warrant transmission failure preserves knowledge closure for two reasons: it implies that clause (ii) of the antecedent is false, and it implies that clauses (i) and (iii) of the antecedent can only be true together when the consequent is true.

matter: it could be that warrant does transmit in zebra-like cases but that claims to warrant do not.²¹

If this shift is taken at face value, the discussion generated from his initial mobilization of the closure/transmission distinction is of considerably more limited interest than it might have seemed. In particular, it will rob the knowledge closure advocate of the Wrightean strategy: she is no longer assured that an agent warranted in believing the premise in a Zebra-like case will end up also being warranted in believing, and so knowing, the conclusion. Perhaps a kind of closure over legitimate knowledge claims is preserved: an agent cannot legitimately claim knowledge of the relevant P without already being in a position to legitimately claim knowledge of Q. But, precisely because of the distinction the later Wright emphasizes, that does not imply knowledge closure, which specifically concerns possession of knowledge itself, not the legitimacy of claims to know.

In order to maintain relevance for the issue of knowledge closure, I will continue to assume that Wright's view pertains to warrant-transmission itself, and so that he advocates the Wrightean strategy in defense of knowledge closure.

VI. Allegiance to WTP

The Wrightean strategy provides an attractive way out for the closure advocate who is uncomfortable with the idea that Bob could learn that it's not a disguised mule by inferring so from his belief that it's a zebra. Certainly intuition favors this resolution: however strong the intuition in favor of WTP might be in general, it is very unintuitive to suggest that Bob could learn that the animal's appearance isn't misleading by inference from a belief grounded in the animal's appearance itself.

Indeed, the price of allegiance to WTP becomes rather exorbitant when one contemplates the number and variety of zebra-style cases (examples of which are constructible for virtually all instances of knowledge).²² A fan of WTP would have to endorse the following inferences as capable of conferring warrant (and so knowledge): from the content of a newspaper report alone to the proposition that that very report is not erroneous as a result of a misprint; from noting that the tank is empty by consultation of the gas

²¹ See, for example, Wright 2012, 471.

²² The construction requires only the possibility of an agent's having the same evidence, or what appears to her to be the same evidence, compatibly with the known proposition P's being false. Take R to describe a scenario in which that is the case. Then a Zebra-style conclusion Q with P as premise can be constructed: $Q = \neg(P \ \& \ R)$. Whether or not the requisite possibility applies to all knowledge, it certainly applies to enough that an alternative not contemplated in the text above—preservation of closure by denial of warrant for P in such cases—amounts to a very thoroughgoing scepticism.

gauge to the proposition that the gauge didn't just break and is stuck on "E" although the tank isn't empty; from one's knowledge that the U.S. president is Barack Obama to the proposition that he didn't suffer a heart attack 5 seconds ago (as a result of which Joe Biden would become president); from one's recollection of where one parked one's car this morning to the proposition that one's car was not stolen from its parking space minutes ago; and so on.

Or, to continue the zebra case: The "Zoo-Testing-R-U's" inspection team, charged with determining that no funny business is afoot at the zoo, peers within every cage and paddock, infers from what they see that the animal therein is not a cleverly disguised something-other-than-what-it-appears-to-be, and comes away satisfied: this zoo is legit. (Indeed, the inspection team is delighted to discover that this is so at every other zoo they visit.)

That price must also be paid with respect to the Moorean inference from "I have hands" to "I am not a handless brain in a vat". While a number of contemporary views are Moorean (or neo-Moorean) in the sense that they endorse knowledge of the latter, very few are willing to claim that the inference from the former to the latter is capable of conferring such knowledge.

Most of these examples are staples of the literature. (In honor of Dretske's having introduced them, call them "Dretske cases".) Closure advocates claim that the conclusions of Dretske cases are known—they must so claim, on pain of scepticism—and offer various compensating mechanisms to accommodate the intuition that they are not. But it is rare for such advocates to claim that knowledge of, and so warrant for, the conclusion can be acquired from the premise in virtue of performance of this particular inference. Neither safety theories nor contextualist accounts, for example, do so.²³

Notice that it would be misguided to insist that we must countenance these inferences as legitimate because they would otherwise constitute exceptions to the axiomatic KC.²⁴ There is, after all, the Wrightean strategy to pursue which, if successful, reconciles the concession that WTP has its limits with KC. But even if that strategy fails, so that the only way to defend KC is to insist on the applicability of WTP to even these cases, such an argument is motivationally backward. It is the intuition encoded by WTP, which WTP raises to the status of a universal principle, which underlies KC in the first place. If it is conceded that, but for our commitment to KC, we would recognize exceptions to WTP then continued endorsement of KC as itself exceptionless is unmotivated as well.

So it seems that the way forward for the closure advocate is to concede that there are limits to the intuition that WTP encodes and to offer an analysis of the cases lying beyond those limits that implies that knowledge is

²³ See, for example, Pritchard 2007, Cohen 1988, and DeRose 1995.

²⁴ Peter Klein nevertheless seems to offer such an argument in, for example, Klein 2004.

closed in those cases as well, as per the Wrightean strategy. I will assume hereafter that advocacy of WTP as exceptionless is simply not a viable option and address what follows to those who agree. Those who disagree, meanwhile, should recognize that allegiance to WTP as underwriting their commitment to KC is much more demanding, and correspondingly more implausible, than is merely allegiance to KC itself.

VII. Explanatory Adequacy

The Wrightean strategy requires:

- (i) that antecedent warrant for the conclusion is a necessary condition of the acquisition of warrant for the premise in the manner in which the latter is acquired;
- (ii) that transmission does fail; and
- (iii) that (i) provides the best explanation for (ii).

For if (i) is not the best such explanation for (ii)—if it is not an adequate explanation and/or there are alternative explanations of transmission failure, ones that do not imply that the conclusion will be warranted but that are at least as plausible as (i)—then we are left without a compelling reason to believe that knowledge is inevitably closed when transmission fails, and so that KC is true.

Condition (iii), therefore, requires:

- (iv) that (i) is an adequate explanation for (ii); and
- (v) that there are no alternative explanations that are at least as plausible.

I will argue that both (iv) and (v) are false. (They are, of course, related: if (i) is not an adequate explanation of (ii), then any alternative adequate explanation of (ii) is the more plausible.²⁵) Since (iv) and (v) are false, so is (iii). But then there's no reason to endorse (i); the only reason to do so is precisely that it provides the explanation for (ii) in the relevant cases.²⁶ But

²⁵ This requires that neither plausibility nor adequacy implies truth: two or more incompatible hypotheses can offer equally plausible, adequate explanations of a phenomenon. This requires a sense of "E explains P" that does not imply that E is true. I assume that such a sense exists.

²⁶ At least, these are certainly Wright's grounds for doing so. One might instead be tempted to suggest that it is strongly intuitive that warrant for the conclusion is presupposed in the relevant cases. But this is not so (especially when it is kept in mind that it is P-warrant that is relevant). See section IV above.

(i) grounds the Wrightean strategy's defense of KC. Without a reason to endorse (i), therefore, advocacy of KC itself is unmotivated.

VIII. First Objection to (iv): Second Warrants

In way of examining (iv), assume that warrant for the conclusion is indeed a precondition of the premise's warrant. Does that fact generate an adequate explanation of transmission failure?

Answering 'no' might well strike the reader as just silly: of course you can't legitimately infer to a conclusion you'd have to be warranted in believing beforehand in order to learn the premise. And indeed, Wright's own writings suggests this response: he repeatedly infers that transmission fails from the assumption that warrant for an argument's premises can proceed only via warrant for its conclusion, without comment.²⁷

To see that the question is not as silly as it seems, recall that I can have multiple warrants for the same proposition. So the fact that I already know, and so possess a warrant for, a proposition does not prevent my acquiring a second warrant for that proposition by inference.²⁸ So it doesn't, in itself, imply that I cannot acquire a second warrant by an inference to that proposition, where the premise's warrant required prior warrant for that same proposition.

Indeed, if the sole reason why warrant does not transmit from premise to conclusion is that warrant for the conclusion is a precondition of the premise's warrant, so that acquisition of the premise's warrant is blocked because one does not yet possess the requisite warrant for the conclusion, then it is hard to see why prior acquisition of warrant for the conclusion by means that are independent of, and so that do not rely on, the argument being mounted would not then allow transmission to succeed.

Again, one could not acquire one's first warrant for the conclusion that way. For that reason one could not acquire knowledge of the conclusion that way either, since knowledge can only be acquired by means of a first warrant.²⁹ As indicated earlier, knowledge, unlike warrant, does not admit of multiple instantiation for the same agent and proposition. But it is not clear why one could not thereby acquire a second (or third. . .) such warrant. So if this were indeed the correct explanation for transmission failure, one

²⁷ See the quote from Wright 2011 in section XV below, for example.

²⁸ Wright is well aware of this, noting that the issue of transmission failure does not merely concern the acquisition of one's first warrant, but also of additional warrants for a proposition one already is warranted in believing. See Wright 2011, 83–84.

²⁹ That is, unless the inference generates a first-time belief for a proposition for which the agent already enjoys a warrant—in which case she acquires both a second warrant and a first belief for a proposition she now knows—or it generates a new ground for a belief already warranted but not grounded in that prior warrant.

would expect transmission to succeed if independent antecedent warrant for the conclusion is secured.

The intuition that transmission fails in the relevant cases, however, is entirely unaffected when such warrant is explicitly provided. Suppose the zoo Bob visits is a shabby tourist-trap petting zoo that is unlikely to be able to afford a zebra, and so which might well resort to a paint-job on a mule.³⁰ Presumably even Dretske would then concede that Bob needs antecedent warrant for the proposition that the animal he is viewing is not disguised (and so not a disguised mule) in order to acquire warrant for the belief that it's a zebra on the basis of its appearance; the disguised-mule possibility is too much of a realistic threat to be discounted.

Fortunately, however, Bob happens to meet the National Zoological Society's zoo-inspection team, which has just examined this particular zoo in order to ensure that no such deceptions have been perpetrated, and who assure him that all is well. (Unlike Zoo-Testing-R-U's, the NZS team employs legitimate, non-bootstrapping methodology, such as DNA analysis.) And so, having glanced inside the paddock, he believes (and presumably knows, this viable threat to such knowledge having been neutralized) that it is indeed a zebra.

"It's a zebra," he asserts; "I know that on the basis of its appearance in conjunction with the fact that it's not disguised (and so not a disguised mule), which I learned from the inspectors. But, now, wait: since it's a zebra, it's not a mule. And since it's not a mule, it's not a mule disguised to look like a zebra. The inspectors were right!"

This is intuitively as bad as ever. While he possesses antecedent independent warrant for the conclusion thanks to his encounter with the zoo inspectors, it remains just as absurd as before to suggest—as suggested by Bob's soliloquy—that he can acquire a second, independent warrant for that proposition in virtue of the inference, one that confirms what he learned from the inspectors.

It is, thus, unclear how citation of the fact that the premise's warrant requires antecedent warrant for the conclusion could deliver an explanation of transmission failure. The obvious explanation along these lines is that the agent can't rely on the argument being mounted in order to acquire the warrant for the conclusion that she needs for the premise's warrant. But then prior, independent provision of warrant for the conclusion should remove that obstacle and allow transmission to proceed. But it doesn't. If the fact (assuming it to be one) that the premise's being warranted requires antecedent possession of warrant for the conclusion is somehow part of the

³⁰ This happened, in fact, in a petting zoo in the Gaza Strip (except that the animal was a donkey). Associated Press, "Gaza Zoo Replaces Zebras with Painted Donkeys", Science on NPR News.com, Dec. 2009. Web. Aug. 11, 2013.

explanation of transmission failure in Dretske cases, the rest of that explanation remains to be provided.

IX. Questions Begged

Of course if the circumstances are such that only a first-time warrant for the conclusion will do, then a warrant for the premise that presupposes warrant for the conclusion in some way will obviously not suffice. If, for example, I am attempting to convince a skeptic that the animal in question is not a disguised mule, and there is no question but that I would have to be warranted in believing that already in order to acquire warrant by perception that it's a zebra (because it's a shabby zoo), then obviously I cannot convince that sceptic by suggesting that it's a zebra and that its being so settles the issue by inference. The charge that I can't legitimately so infer because doing so presupposes what I am trying to prove, and therefore begs the question, is entirely in order.

But in the shabby zoo case Bob is not attempting to convince anyone—neither himself nor anyone else—of something that person doubts, or is open-minded about, before the argument's presentation. He fully recognizes that, in light of the zoo's shabbiness, he must have a prior independent warrant for the proposition that the animal is not a disguised mule (which is fortunately available thanks to the inspectors) before he can legitimately determine that the animal is a zebra on the basis of its appearance. With that warrant in hand—and so with the question settled in the proposition's favor—he then proceeds to attempt to acquire a second warrant by performing the inference. And he cheerfully concedes that it is only a second warrant that can be acquired in such a fashion (or so he thinks). So understood, this is not a case in which Bob is illegitimately assuming the very proposition that his sceptical or open-minded interlocutor takes to be initially at issue, and whose truth he is attempting to establish to the satisfaction of both.

This is a contextual matter: there are some people in some circumstances (which may include Bob himself) who Bob cannot legitimately convince by appeal to this argument, in virtue of their initially sceptical or open-minded belief-states. To offer the argument to such people is to beg the question. But it is an open question whether such an argument can deliver a warrant to those who are not initially sceptical or open-minded. That an argument would beg the question in one context does not demonstrate that it does so in all.

X. Argument Structure and Inferential Path

Suppose Bob is asked why he believes that the animal is not a disguised mule. He points out that it's a zebra and zebras are not mules (disguised or not). He is then asked why he believes that it's a zebra; he testily responds that it looks like one. Finally, when pressed to consider the possibility that

it only looks like a zebra because it's actually a disguised mule, he explodes, "I just told you: it's a zebra. Zebras aren't mules!"

This does not exhibit merely circularity of argument structure (because the conclusion, or warrant for it, is in some way presupposed by the argument). It also exhibits circularity of inferential path. Bob is appealing to the conclusion, along with the warrant he takes himself to have earned by that very argument, to answer a question concerning why he is warranted in believing the premise.

This highlights an ambiguity in the charge that the arguer assumes the conclusion in her inference to that very conclusion. If the reasoner relies, not just upon the proposition expressed by the conclusion, but upon it as the conclusion of that argument—as warranted by that very argument—then the path the reasoning takes is circular. But the reasoner can instead take herself to have antecedently established, by independent means, the proposition expressed by the conclusion in the process of establishing warrant for the premises, and so without engaging in circular reasoning.

The Wrightean strategy's explanation of transmission failure—that antecedent warrant for the conclusion is presupposed by the warrant for the premise—succeeds when the inferential path is circular. If Bob relies on warrant for the conclusion as provided by the very argument under construction in support of that argument's own premises, transmission indeed fails. Warrant for the conclusion cannot be acquired in this fashion until the argument is complete; but it can't be completed until he already possesses warrant for the conclusion. There is no way to get started.

But while Bob's argument in the shabby zoo case does exhibit circularity of argument structure, his reasoning does not trace a circular inferential path susceptible to that explanation. While he does rely on the proposition that it's not a disguised mule while deciding, on the basis of its appearance, that it's a zebra, he explicitly does not rely on that proposition as warranted by the very argument he is attempting to mount. That proposition is supported by the independent evidence provided by the zoo inspectors. With that in hand, he then arrives at the premise that it's a zebra, and proceeds from there to the conclusion that it's not a disguised mule. The conclusion indeed expresses the same proposition as that upon which he initially relied. Nevertheless, he is not employing that proposition qua conclusion of that argument, that is, as putatively warranted by the inference the argument represents. So it remains a mystery why transmission would nevertheless fail, given only these explanatory resources.

XI. Second Objection to (iv): Mismatched Strength of Intuition

Nor does the strength of the intuition that transmission fails vary with the sense that independent antecedent warrant for the conclusion is required.

However plausible it is that such warrant is required in the shabby zoo case, as Dretske pointed out it is somewhat less intuitively plausible that Bob either does or must possess such a warrant when visiting a more typical, reputable zoo. There is a continuum of cases ranging from those in which it is intuitively quite plausible that such antecedent warrant is required to those where it is somewhat less plausible.³¹

In way of underscoring this, note that most are comfortable with the idea that young children can know that it's a zebra on the basis of its appearance. But not only do they surely not contemplate the enormous number of ways in which some form of deception might be taking place, they may even lack the conceptual resources to do so (with respect to those deceptions involving holographic projection, say). It is at least odd to suggest that, in run-of-the-mill zoos, children require, let alone possess, warrant for the belief that no holographic projections are taking place in the zebra paddock. Not so, however, in the shabby zoo: the threat of disguising (albeit by more prosaic means) is too proximate for a child to learn, by looking at the animal, that it's a zebra. The child, no less than the adult, intuitively needs the wherewithal to eliminate that threat, whether or not she is in a position to do so.³²

The intuition that transmission fails is, however, equally strong no matter where the case falls on the intuitive-plausibility scale. That is, whether the zoo is shabby so that Bob must, intuitively, be warranted in believing that no animal-disguising has taken place, or it is reputable and the need for such antecedent warrant is considerably less obvious, it remains intuitively highly objectionable—equally highly objectionable—to suggest that adult or child can acquire a warrant, first or second, by inference from “it's a zebra” to “it's not a disguised mule”. This is not what one would expect if the explanation for transmission failure essentially appealed to the need for such antecedent warrant. The degree to which the inference is intuitively objectionable should vary with the degree to which antecedent warrant for the conclusion is intuitively required.

This remains so even if, in fact, antecedent warrant is required in all cases (sadly so for the children). Those who insist that this is so will need an error theory to explain the intuitive misperception that the need for antecedent warrant is stronger in some cases than in others. But whether or not

³¹ Recall again that the warrant in question is P-warrant—that which takes true belief into knowledge—rather than merely that the agent has some background reason to believe that animal disguisings are unlikely. There is surely variation in the intuitive force that P-warrant—knowledge, if the agent believes it—is needed.

³² We zoo-visiting adults are, moreover, in the same position as are our less conceptually sophisticated children with respect to some ways in which such deception can be perpetrated. See section XIII below.

the intuition is veridical or illusory, it exists. And since it exists, the degree to which the inference is intuitively objectionable should vary with it.

The upshot is that citation of the fact, assuming it is one, that warrant for the premise presupposes warrant for the conclusion does not adequately explain transmission failure. If it did, transmission should succeed, albeit inevitably to a second warrant, so long as (i) the argument is not being offered to someone initially sceptical or open-minded about the conclusion and so does not beg the question, and (ii) antecedent warrant for the conclusion is delivered prior to and independently of the argument being mounted (so that the result does not exhibit circularity of inferential path).³³ But it doesn't seem to; transmission is as intuitively unacceptable as ever after these conditions are met. We would, moreover, expect our intuitions concerning whether transmission fails to vary in strength with our intuitions concerning whether such antecedent warrant is required; but it doesn't.

XII. Wright vs. Wright

This result is somewhat surprisingly bolstered by Wright's own recent work on the issue. He provides an example of an argument wherein the conclusion must be warranted beforehand in order for the premise to acquire its warrant, but which Wright views as at least arguably succeeding in transmitting warrant.

[S]uppose I am trying to count a rather mobile, smallish flock of sheep in a pen. And suppose I am told that there are fewer than twenty and, on counting as carefully as I can, get the result eighteen. I might reasonably take that result as corroborating what I have been told. But I might also reasonably think that, had I lacked any independent information about the number of the flock, I could not in the circumstances have trusted in the result of the count sufficiently to conclude that their number was fewer than twenty.³⁴

I infer from my having counted 18 that there are indeed less than 20 sheep. But I could only reasonably trust that count—my counting could only

³³ It is disputable whether (i)'s being false really does deliver transmission failure, rather than simply implying the unavailability for dialectical purposes of an argument that might, nevertheless, transmit. This latter is the view of James Pryor, who claims that transmission succeeds in the Moorean inference from "I have hands" to "there is an external world", but which argument is nevertheless dialectically unavailable against a sceptic. See, for example, Pryor 2004. In way of reconciliation, Martin Davies 2009 suggests that these are distinct forms of transmission failure, corresponding to distinct epistemic projects. See also Coliva 2012. (This relevance of this literature for the present topic—knowledge closure—is however limited by the problem of the ambiguity of warrant discussed in section IV above.)

³⁴ Wright 2011, 83.

provide warrant for belief that there are 18 sheep—against the background of my having been told, and so already being warranted in believing, that there are less than 20. And yet, Wright concedes that transmission might succeed here. He introduces the example, indeed, as a possible case in which transmission succeeds in generating, not a first-time warrant, but a second warrant for a proposition already, and of necessity given the circumstances, warranted.

If we assume that Wright's conservative views with respect to perceptual warrant are correct, then another such case—in which transmission appears to succeed notwithstanding the need for antecedent warrant for the conclusion—becomes plausible. According to Wright (vs. Pryor) we need antecedent warrant for the belief that perception is reliable—that we are not brains in vats, victims of an evil genius, etc.—before we can legitimately take ourselves to acquire warrant on the basis of perceptual experience. Conceding that such a warrant cannot itself be delivered by such experience, Wright suggests we construe it as a default entitlement to believe (or at least accept), by courtesy of which entitlement we may proceed to acquire warrant on the basis of perception for beliefs about the surrounding world.³⁵

Suppose that is correct. Thanks to that entitlement's being in place, we appear to have acquired, by scientific and so ultimately perceptual means, a wealth of information concerning the anatomy, neurophysiology, physics, etc. of our perceptual systems and the sophisticated way in which they generate information concerning our perceptually accessible physical environment. Putting that information together appears to deliver the result that our perceptual system does indeed reliably indicate the character of our physical surroundings, at least in the right environmental circumstances and within certain specifiable limits, and indeed to reveal the physical processes that underlie that reliability.³⁶

We then seem to have acquired a second, empirical warrant for a proposition—that perception is, at least within certain limits, reliable—which was initially warranted as an entitlement. This second warrant, moreover, requires that the entitlement-warrant initially be in place (on Wright's view). For if it weren't, then I would not be warranted in believing any perceptually based proposition, and so not warranted in believing the empirical claims that ground my second empirical warrant for this same proposition.³⁷

³⁵ See Wright 2004.

³⁶ This is a far from trivial result and one which is indeed limited. For example, we've long known, by empirical means, that the apparent retrograde motion of the planets as discerned from our intuitively fixed vantage point on the Earth's surface is an unreliable indicator of those bodies' actual movements.

³⁷ To be clear, I'm not endorsing Wright's view here. I merely point out that, having taken that view on board, it would seem reasonable to countenance a second, empirical warrant for the same proposition initially warranted as an entitlement.

Notice that, while this is a case wherein warrant for the conclusion is presupposed by the argument through which the second warrant was acquired (or so we are assuming with Wright), it is not a case in which the reasoning itself traces a circular path. It is not the proposition expressed by the conclusion along with the warrant for it delivered by that very argument that grounds that same proposition when it appears as a premise. In the latter guise the warrant is a default entitlement, whereas in the former guise it is an empirically grounded warrant.

Notice also that this is not (or need not be) a case in which the question is begged. Of course, if a sceptic refuses to recognize any such thing as a default entitlement for the claim that perception is reliable, then an argument that delivers an empirical warrant for that claim but whose premises presuppose a default entitlement to the same claim will not move him. Wright does offer considerations in favor of the existence of such warrants. But a primary such consideration is the suggestion that this is the only way to avoid scepticism, a consideration that will hardly influence the sceptic. Nevertheless, it is not obvious that such an argument could not provide those who do not start out in a sceptical frame of mind with a second warrant for belief that perception is reliable, given Wright's own views.

So transmission does not obviously inevitably fail whenever warrant for the premise presupposes warrant for the conclusion. One of Wright's own examples suggests otherwise, as does a combination of his views on perceptual entitlement warrants together with the implications of the information available by courtesy of such warrants. He can then hardly claim that transmission failure is generated whenever the conclusion must be antecedently warranted in order for the premise to enjoy its warrant, so that the latter requirement explains the former failure.

XIII. Third Objection to (iv): The Demands of the Wrightean Strategy

A final objection to the Wrightean strategy, unrelated to its status as an explanatory hypothesis, concerns the demands it makes on the would-be knower. According to the Wrightean strategy transmission fails in Dretske cases because warrant for the conclusion is a precondition of the premise's warrant. Note that the latter must be true whether or not the agent attempts the inference: that it is true explains why transmission would fail were the agent to attempt it, whether or not she does so in fact. So the strategy requires that the agent be antecedently warranted—P-warranted—in believing the conclusion of every constructible Dretske case for every known proposition.

But there are innumerable such cases for every knowledge ascription. Not only must the animal not be disguised; there must also be no intervening screen upon which a zebra-image is projected, Bob must not be

hallucinating, light must travel a straight path from animal to eye, certain complex neurological conditions required for visual processing must be in place, and so on. Some of these—the neurological conditions, for example—will be beyond layperson Bob’s ken. There is (presumably) a complex proposition describing a neurological state that, if realized, would both prevent the processing of visual information and yet eventuate in the same experience. The negation of that proposition is therefore a condition of Bob’s knowing that it’s a zebra. But Bob is in no position to understand that proposition, let alone recognize the threat that such a state would present to his zebra-knowledge.

The Wrightean strategy requires, nevertheless, that Bob be warranted—P-warranted—in believing that this proposition is true. That’s an awful lot to expect of Bob. More generally, the acquisition of much, if not all, of our knowledge requires the realization of any number of background conditions, some of which might be recognizable as such by the agent but many of which will not be. Worse, this iterates: knowledge that, and so warrant for belief that, a particular such condition holds will itself be achievable only against the background of further conditions.

Perhaps Bob enjoys a default entitlement to some of the more general such conditions. But that’s surely not so for a great many others (such as those mentioned two paragraphs above); warrant for those presumably requires something in the way of empirical evidence. This is surely so for the denial of the neurological-state proposition. So Bob must possess a warrant, gleaned on the basis of empirical evidence, for denial of a proposition he is no position to understand, let alone recognize as describing a threat to his knowledge; and similarly for all other non-entitlement conditions (and conditions of those conditions). That way lies skepticism.

XIV. Warrant-Circularity

We have been pursuing the question whether assumption (iv) of the Wrightean defense of KC is true: assuming that knowledge of the conclusion is indeed presupposed by the warrant enjoyed by the premise, does that fact generate an adequate explanation of transmission failure? We will now pursue the question corresponding to assumption (v): is that the only explanation available? Assuming that there are indeed instances of transmission failure, and so exceptions to WTP, our having answered ‘no’ to the first question implies the same answer to the second: the explanation for transmission failure must lie elsewhere.³⁸ What then explains transmission failure in Dretske cases, if it is not a matter of the conclusion’s having to be warranted or known beforehand?

³⁸ Assuming, of course, that some such explanation exists.

Two characteristics are consistently manifested by such cases; their interaction, I suggest, provides the explanation. The first characteristic, with which the remainder of this section is concerned, is that the conclusion is an *enabling condition* for the premise's warrant: it must be true in order for the agent to acquire warrant for the premise in the manner in which that warrant is actually acquired.³⁹ Whether or not Bob must be warranted in believing that it's not a disguised mule, for example, in order to acquire a warrant for his belief that it's a zebra on the basis of the animal's appearance, it certainly must not in fact be disguised (and so not a disguised mule).⁴⁰

This is so no matter what belief is thereby warranted. Determining an animal's species membership on the basis of its appearance requires that its appearance be a genetic manifestation of species membership in the usual, natural fashion, no matter what sort of animal it is.⁴¹

As a result, Bob's warrant could be blocked because this requirement fails even if it is in fact a zebra. A likely story: the zoo was sold what it took to be a mule, but which was in fact a zebra with a skin disease that erased its stripes. The zoo authorities then painted it to look like a zebra and put it in the zebra paddock in order to fool their visitors. And then along came Bob... This is a Gettierized version of the disguised-mule scenario: Bob has acquired a justified, true belief that it's a zebra, but he doesn't know it. As per section IV above, Gettierized beliefs are not P-warranted.

³⁹ I will discuss the second characteristic in the next section.

⁴⁰ One might object that this leads to the result that P-warrant for A implies that A is true, notwithstanding my earlier professed neutrality on the issue in section II above. Enabling conditions arguably concern either the appearances themselves—it must look like a zebra—or they eliminate ways in which that appearance could be misleading (it's not disguised, Bob's not hallucinating, etc.). But if the appearances are as though A and those appearances are not misleading, that seems to imply A. I am inclined to bite the bullet and draw the conclusion that P-warrant does imply truth. After all, a P-warrant is, as it were, a way in which one can come to know that A, a path that leads one to its truth. So understood, it'd be no surprise that such a way or path must in fact arrive in order to count as such. But an alternative is to weaken the condition to "if A is true then the condition is realized"—"if it is a zebra then it's not disguised", for example—so that the condition is realized whenever A is true. There is no room to evaluate this alternative here. However, so far as I can see, no harm comes to the arguments mounted in this paper if the alternative is substituted throughout. So I will continue to use the stronger versions. (Doing so is, at least, clearly in line with Wright's presentation of these conditions; he introduces no such qualifications when discussing them.) Thanks to James Pryor for both bringing my attention to the issue and presenting the alternative.

⁴¹ This is another point at which it matters that it is P-warrant rather than W-warrant in play. I might well be W-warranted for "that's a zebra" when it's disguised (especially if W-warrant is internalist-style justification). But I can't be P-warranted: I still won't know that it's a zebra, even if it is. (How could it be a disguised zebra? Read on.)

As this highlights, the feature of the disguised-mule scenario that is incompatible with the warrant for “it’s a zebra” is distinct from that which ensures that “it’s not a disguised mule” follows from that proposition. Knowledge being factive, it can’t be a mule, and so not a disguised mule, compatibly with Bob’s knowing that it’s a zebra. But there is a second incompatibility: Bob can’t acquire that knowledge, in the manner in which he acquired it, compatibly with it’s being disguised, and so a disguised mule, irrespective of the facticity of knowledge.

Moreover, it is the incompatibility of the animal’s being disguised, rather than its being a mule, that is correlated with the intuition that transmission fails. There’s nothing intuitively wrong with Bob’s inferring “it’s not a mule” from “it’s a zebra”; it’s the possibility of disguise that causes problems. The inference relation between premise and conclusion, however crucial to the question whether KC holds, directs attention away from the crucial feature of the relation between knowledge of the premise and truth of the conclusion responsible for transmission failure.

That the animal must not *in fact* be disguised (and so not a disguised mule) in order for Bob to be warranted in believing that it’s a zebra does not, on its own, imply that Bob must be *warranted* in believing that it’s not disguised. The dependence relation involved—call it warrant-circularity—is more akin to rule-circularity than premise-circularity in this respect.

An argument is premise-circular when the conclusion is itself a premise of that very argument, whereas the conclusion of a rule-circular argument affirms the reliability of the rule of inference that argument instantiates.⁴² Warrant for a premise-circular conclusion is obviously a condition of warrant for the premises, since the conclusion is itself one of the premises. This is less obvious for rule-circular conclusions: perhaps the rule need only be reliable in fact. The nature of rule-circularity itself, unlike that of premise-circularity, does not imply that prior warrant for the conclusion is required; there is room for debate.⁴³

There is similar room for debate with respect to warrant-circularity. That an inference is warrant-circular does not, in virtue of that characterization alone, imply that the premise’s being warranted requires that the conclusion be warranted rather than just true.

Notwithstanding his views on the matter, Wright’s most recent “template” for transmission failure also leaves this issue open. One cannot, he suggests, acquire warrant by inference to a conclusion that describes a

⁴² Some advocates of rule-circular reasoning might have a different characteristic than reliability of the rule (and so a different sort of conclusion in rule-circular arguments) in mind, although reliability is by far the most commonly cited property. It won’t matter for the issue at hand.

⁴³ Advocates of rule-circular arguments include Braithwaite 1953 and Psillos 2004.

“presupposition of the cognitive project of obtaining a warrant for [the premise] in the relevant fashion [being that in which the agent is in fact attempting to acquire such a warrant]” (Wright 2011, 93), where a “presupposition of a cognitive project” is “any condition P such that to doubt P (in advance of executing the project) would rationally commit one to doubting the significance, or competence of the project, irrespective of its outcome” (Wright 2011, 93).⁴⁴

To doubt that it’s not a disguised mule, for example—that is, to suspect that it is a disguised mule, or at least to be unconvinced that it isn’t—would rationally commit one to doubting the significance of the project of appealing to the animal’s appearance in way of acquiring warrant for belief in its species (regardless of the species to which its appearance might suggest it belongs).

The fact that doubting the conclusion rationally requires doubting the significance of the cognitive project by which warrant for the premise was acquired implies only the rational incoherence of combining such a doubt with endorsement of that warrant. But it does not imply that one must be positively warranted in believing the conclusion in order for the premise’s warrant to go through. The simplest explanation for why doubt about the conclusion undermines the project, after all, is that one recognizes that the conclusion must be true in order for the project to succeed: if one does so and yet doubts the conclusion, that should obviously undermine one’s confidence in the project’s success to some extent.

So it is compatible with the conclusion’s being a presupposition of the cognitive project involved in delivering a warrant for the premise that the agent not be warranted in believing that conclusion. To claim, with Wright, that the agent nevertheless must be so warranted, therefore, requires argument. It will not fall out of the characterization of warrant, and of presupposition, by itself.

Wright’s concept of a presupposition is, admittedly, richer than the concept of an enabling condition. The fact that doubt about a presupposition rationally requires doubt about the project of obtaining a warrant itself requires that the agent recognize that the presupposition must be in place in order for the project to succeed. But reflection on the neurological-state

⁴⁴ The qualifier “irrespective of its outcome” is needed because, as Wright points out, a prior opinion against P might undermine the project; P would then itself count as a presupposition. Antecedent conviction that it’s not a zebra (because one seems, incorrectly, to remember that the zebra paddock was on the other side of the zoo) might count against the project of determining what the animal is by looking. But its being a zebra is not itself a presupposition of that project.

In his 2012, Wright refers to presuppositions as “authenticity-conditions”. I retain “presupposition” here in line with his 2011, in which the template discussed in section XV below is explicitly presented.

example from the previous section indicates that this is too narrow. Even if Bob could comprehend the proposition, this does not ensure that he recognizes that it must be in place in order for him to learn that there's a zebra in the paddock by looking. So he will not be rationally required to doubt that project in virtue of doubting that proposition, simply because he fails to recognize its relevance for that project. Transmission to that proposition is, however, intuitively as problematic as it is to propositions (like "it's not a disguised mule") whose relevance Bob is in a better position to recognize. The concept of an enabling condition, whose realization does not require that it be identified as such by the agent, is therefore the preferable one to wield in this context. None of this matters, however, for the present point: neither the concept of a presupposition nor the concept of an enabling condition implies that the agent is antecedently warranted in believing that presupposition/condition to hold in order for the project involved to succeed.

XV. Indistinguishability

The second characteristic exhibited by Dretske cases is captured in Wright's new template for transmission failure as well. Here is the full statement of the template.

Revised Template

Where P entails Q, a rational claim to warrant for P is not transmissible to Q if there is some proposition R such that:

- (i) The process/state of accomplishing the relevant putative warrant for P is subjectively compatible with R's holding: things could be with one in all respects exactly as they subjectively are yet R be true.
- (ii) R is incompatible (not necessarily with P but) with some presupposition of the cognitive project of obtaining a warrant for P in the relevant fashion, and
- (iii) Not-Q entails R.⁴⁵

⁴⁵ Wright 2011, 93. It is titled the "revised template" because it is designed to unify his earlier "disjunctive" and "information-dependent" templates. See, for example, Wright 2000 for the former and Wright 2002 for the latter. Note that this template concerns the transmission of a "rational claim to warrant", and so not necessarily warrant itself. As per section V above, in the interest of maintaining the relevance of his views to the issue of knowledge closure I will nevertheless treat this as a proposed template for P-warrant transmission failure.

I have replaced "A", "B" and "C" in the original with "P", "Q" and "R" for consistency's sake here and below.

In the zebra case, P is “it’s a zebra”, Q is “it’s not a disguised mule”, and R—“It is a disguised mule”—just is not-Q.⁴⁶

The second characteristic appears as condition (i): the “subjective compatibility” of R’s holding with “the process/state of accomplishing the relevant putative warrant for P”. It could be a disguised mule, for example, compatibly with Bob’s having the same subjective experience he actually has in the presence of the real zebra.

Why does Wright’s revised template include this feature? If the key to transmission failure is that the conclusion represents a presupposition of the cognitive project involved in delivering a warrant for A, why should it matter that, in addition, failure of that presupposition to hold is subjectively compatible with the process/state involved in accomplishing that warrant?

Wright’s answer is as follows.

Let us run through the zebra case. R [“those animals are mules cleverly disguised to look just like zebras”] is incompatible with P and is the negation of Q. So conditions (i) and (ii) are met. I look at the animals in the zoo enclosure and, so it seems to me, perceive that they are zebras. But the subjective state involved in such a perception is shared with a perception of mules, artfully disguised to look exactly like zebras (condition (i)). So if I am rationally to claim to be perceiving zebras, and to possess warrant for P on that basis, I better be in a position to discount the mule hypothesis. (Otherwise the suggestion that it only seems to me that I am perceiving zebras is open.) But that is the same as to say that I better be in a position to affirm that the animals in question are not cleverly disguised mules—precisely Q. So that is something that I need to be in position to claim on independent grounds prior to claiming perceptual warrant for P, and on which my claim to warrant for P therefore rests. It is accordingly not something for which I can acquire a claim to warrant by the argument itself.⁴⁷

There are, however, three problems with this, at least insofar as it might be employed to demonstrate that transmission fails because prior warrant for “it’s not a disguised mule” is required. First, even assuming that I must be “in a position to discount” the disguised mule hypothesis, this does not obviously require that I be warranted in disbelieving it, let alone that I am P-warranted in doing so. Pryor, for example, would claim that being in such a position requires, at least in some cases, no more than not having evidence counting in favor of it. To employ this argument against such a position therefore begs the question. And others might concede that being in such a position requires some background evidence against the disguised

⁴⁶ Alternatively one could, compatibly with the template, treat R as “it’s disguised”.

⁴⁷ Wright 2011, 90.

mule hypothesis, without requiring enough such evidence to count as knowing it.

Second, even if I must be warranted in believing that it's not a disguised mule in order to legitimately claim warrant for "it's a zebra", it is compatible with this that it must simply be true in order for me to in fact possess such a warrant.⁴⁸

Third, even if I must be warranted in believing that it's not a disguised mule in order to be warranted in believing that it's a zebra, the claim of the last sentence of the quotation—that this immediately implies and so explains transmission failure—is false. That was the lesson of sections VIII–XIII above.

In sum, Wright's new template incorporates what I believe to be the two crucial characteristics required for the explanation of transmission failure: warrant-circularity and indistinguishability.⁴⁹ He has not, however, correctly identified how they contribute to, and so explain, that failure.

XVI. Procedural Self-Defeat

Before attempting such an explanation, I would like to briefly consider a non-inferential example.⁵⁰ On your way to an unfamiliar town, you stop at a gas station just off the highway and ask the attendant the way downtown, who replies, pointing, "it's that way" (D). You have no reason to suspect him of deception and some reason to think that he knows D (the lack of adequate signage suggests that he has been asked this before). And indeed he does know D, and answers truthfully. You presumably now know D as well, in virtue of his testimony. But you then wonder, for no particular reason, whether he is trying to mislead you. So you ask, "Are you trying to mislead me?" He answers, "Oh, no, not at all".

The attendant presumably knows whether he is trying to mislead you at least as well as he knows D. And he answers your question, despite its oddity, as honestly as he did your first question. So his response to this second question is just as knowledgeable and honestly delivered as his response to the first. And yet there is intuitively something seriously misguided in the attempt to acquire *this* bit of information by testimony. Why?

In light of reasoning analogous to that in sections VIII–XI above, the answer isn't "because you need to know that he is not trying to mislead

⁴⁸ Wright is aware of this; as indicated earlier, he explicitly shifts topic from warrant-transmission to transmission of legitimate claims to warrant. Assuming the topic has not shifted, so that Wright's argument is construed as pertaining to the transmission of warrant itself, the objection stands.

⁴⁹ The latter characteristic, however, requires some finessing; see below.

⁵⁰ Or, at least, an arguably non-inferential example. The epistemological grounding and status of testimony is of course a controversial matter.

you in order to learn anything from him by testimony”. It is not obvious that you need to know this in order to learn the way downtown from him; that’s a controversial matter. Perhaps we need (and have) some background reason to believe that people won’t typically attempt to mislead their interlocutors, just as we might need (and have) some background reason to believe that zoos don’t typically disguise their animals. But the former is a long way from saying that you know that this particular person is not misleading you now, just as the latter is a long way from saying that Bob knows that this particular animal in this particular zoo is not disguised.

Nor does the strength of the intuition that you need to know that he is not trying to mislead you vary with the intuition that the question “are you trying to mislead me?” is self-defeating. (The latter intuition is present, for example, no matter whether the attendant wears a suspicious smirk on his face or is, by all appearances, ingenuous.) And even if you did somehow acquire information that he is not trying to mislead you by independent means (a lie-detector built into the Google-glasses you’re wearing, say), to view his answer as itself generating an additional warrant for that claim, corroborating the verdict of your glasses, remains intuitively misguided.

Solicitation of the attendant’s testimony in way of determining whether he is trying to mislead you constitutes a cognitive project in Wright’s sense. Such a project is “defined by a pair: a question, and a procedure one might competently execute in order to answer it”. (In the present case, the pair is “is the attendant trying to mislead me?” and “asking the attendant”.) Generally speaking, a cognitive project is in the service of improving one’s epistemic standing with respect to whatever answer it delivers: one acquires a reason, justification, warrant, knowledge, etc. in favor of that answer that one did not have before executing the project. In line with the ultimate topic of this paper—knowledge closure—I will focus on projects whose purpose is to deliver a P-warrant.

For simplicity’s sake, suppose the relevant questions admit only “yes” or “no” answers (as does “is the attendant trying to mislead me?”, “is it a disguised mule?”, etc.).⁵¹ Call the project’s *target* the proposition affirmed or denied: “the attendant is trying to mislead me”, “it’s a disguised mule”, etc. Finally, call the project’s answer with respect to its target the *outcome* of that project. A consequence of the simplifying assumption is that the outcome, assuming there to be one, affirms either that the target is true or that it is false. If the project succeeds, the agent acquires a warrant to believe the outcome that she did not previously possess.

⁵¹ Of course not all questions are like this; “how many cats are there in the room?” admits of any number of numerical answers. Nothing that I can see of relevance hinges on this in the discussion below.

Such projects succeed when they do in virtue of the realization of a variety of enabling conditions. These are, roughly speaking, the same conditions that Wright refers to as presuppositions (or, later, as authenticity conditions).⁵² One can only learn something by testimony if the testifier is not trying to mislead you; one can only learn that it's a zebra by visual inspection if the animal is sufficiently illuminated; one can only determine the length of a board by means of a tape measure if the measure's markings are well-calibrated; and so on.⁵³ The agent may be aware that the project she has undertaken requires the holding of a particular such condition. But her being so aware is not required in order for it to count as such. It constitutes such a condition, required for the project's success, regardless.⁵⁴

Some such conditions are *opaque*: the project will deliver the same outcome as that actually delivered when they are not realized. Bob looks into the paddock, sees what looks like a zebra, and acquires the belief that it is a zebra on the basis of its so appearing. In order for Bob's animal identification project to succeed, the animal must not be disguised to look like a zebra.⁵⁵ This condition's being false, however—the animal is so disguised—implies that it will still look just like a zebra to Bob. Given that the animal-identification project Bob has undertaken precisely involves relying on its appearance in way of determining whether it is a zebra, implementation of that project when that condition is false implies that the project's outcome will be “that's a zebra”, which is the actual outcome.

Not all enabling conditions are opaque. Bob's animal identification project requires sufficient illumination in order to see the animal. If this condition did not hold—because it's pitch black—the outcome would certainly be different; indeed, it presumably will deliver no outcome at all. This is so no matter what the actual outcome is.

In the zebra case the opacity of the condition that it not be disguised to look like a zebra is a function of the actual outcome, namely, “it's a zebra”. If the outcome is instead “it's a mule”—because it is an undisguised mule, say, or because it's a zebra disguised to look like a mule—then “it's not disguised to look like a zebra”, although still an enabling condition, is not opaque. That condition's falsehood still implies that it looks like a zebra, which would generate the outcome “it's a zebra” by the animal-identification

⁵² See footnote 44 above.

⁵³ I doubt that the concept of a condition-less project is coherent. But even if it is, surely all those undertaken by humans are not; any such will require at least appropriate background psychological capacities.

⁵⁴ See sections XIII and XIV.

⁵⁵ As per the discussion of the diseased-zebra example in section 14, it doesn't matter whether the animal is in fact a mule, zebra, or something else entirely. It's the disguise that matters.

project. That outcome is, in this case, distinct from the actual outcome, “it’s a mule”.

The opacity of some conditions, however, are not outcome-relative in this way. “The animal is not cleverly disguised to look the way it does in fact look” is also a (closely related) condition of the animal-identification project. If it is false—it is cleverly disguised to look the way it does look—then the animal-identification project, relying on how it looks, will generate the outcome that it is the way it looks in fact (whatever that is), which will inevitably be the actual outcome of that project. Call these outcome-independent (vs. outcome-relative) opaque conditions.

That a project has opaque conditions (of either variety) does not in itself imply that it fails. Most—plausibly, all—cognitive projects that humans undertake will have such conditions, as a result of which those projects are defeasible (by subsequent indication that the condition does fail). If the presence of an opaque condition alone implies that the project relying on it is unsuccessful, the result will be rampant skepticism.

Nor does the fact that a condition is opaque with respect to a particular project imply that the agent does not possess a warrant to believe that that condition is realized. She could have acquired such a warrant by project-independent means. Or, perhaps, some warrants can be had notwithstanding the opacity of the condition they warrant, even without access to project-independent sources.⁵⁶

Nor does the fact that a project has an opaque condition imply the inevitability of the outcome of the project for which it is a condition. The opacity of the condition that the animal not be disguised to look like a zebra does not in itself imply that the animal will in fact look like a zebra, so that the outcome “it’s a zebra” will result. If all goes well—the opaque condition holds, as do the other enabling conditions—it will look like a zebra only if it is one.

If a condition’s opacity *did* imply the inevitability of the outcome, however, that would invalidate the project. Suppose that it follows from the nature of the procedure employed by a project that, if the outcome it delivers is false, that procedure will nevertheless still generate that same outcome. Such a project obviously delivers that outcome without regard to the outcome’s truth-value. It can then hardly constitute a method of inquiry in the service of determining whether that outcome is true.

In way of resolving his dispute with his heliocentrist opponents, for example, the geocentrist Tycho Brahe points to the setting sun. “Of course the sun is revolving around the earth”, he exclaims; “you can *see* the disc of the sun sinking down past the horizon on its journey around the earth!”

⁵⁶ I am not, however, suggesting that all—or indeed any—such conditions do require warrants.

If, however, the earth were rotating in relation to a stationary sun as the heliocentrists claim, that apparent motion would be precisely as it is.⁵⁷ The method Brahe employs—appeal to that apparent motion of the sun as indicative of its actual motion—will therefore still generate the answer he favors. It is for that reason a pointless exercise rather than a legitimate cognitive project. No improvement in the epistemic standing of the view Brahe favors—no new warrant—therefore results from that procedure’s employment.

Such projects are self-defeating. There are other ways in which a cognitive project could be self-defeating. The procedure it invokes could be, for whatever reason, impossible to realize.⁵⁸ Or, the procedure might deliver a response that encodes the correct answer, but does so in a manner that is, for whatever reason, inevitably inscrutable to the inquirers who hope to utilize it. The form of self-defeat described two paragraphs above is only one among many. Call it *procedural* self-defeat, in light of the fact that the self-defeat is a consequence of the structure of the procedure involved.

The gas station attendant case manifests procedural self-defeat. If the target proposition—“the attendant is not attempting to mislead me”—is false, the attendant will still affirm that proposition. (If he were to confess that he is attempting to mislead you, that very confession brings to an end his attempt to mislead you.) Since the procedure consists in asking him and treating his response as accurate, the outcome remains the same. So the procedure employed to investigate the target is such that it will still deliver the outcome that it does deliver even if that outcome is false. The project is therefore a failure, as are all procedurally self-defeating projects.

That this is so in this case is a consequence of the fact that the target expresses an opaque condition of the very project being mobilized to evaluate it. It is undoubtedly a condition of that project: the attendant must in fact not be attempting to mislead you in order to learn that he is not attempting to mislead you from his saying that he is not attempting to mislead you. And it is opaque: if it is false—he is trying to mislead you—he will assert that he is not trying to mislead you, which is what he does assert. It is because the outcome expresses an opaque condition of the project involved in evaluating it that it is procedurally self-defeating.

This will be so whenever the outcome of the project specifies an opaque condition of its own successful execution. Since the condition is opaque, it follows from the condition’s falsehood that the project will deliver the same

⁵⁷ In this way the heliocentric account explains the apparent motion of the sun across the sky as a kind of visual illusion, akin to the perception of a train on an adjoining track seeming to move backward when one’s own train is in fact imperceptibly moving forward. The example is from Hanson 1958.

⁵⁸ There are different grades of such self-defeat depending on the modality involved in the claim that the method’s realization is impossible.

outcome that it in fact delivers. But that condition is also the project's outcome. So it follows from the *outcome's* falsehood that the project will deliver that same outcome. Any such project is therefore procedurally self-defeating.

The testimony scenario presents a doomed cognitive project, then, because (a) its target describes an enabling condition of the very project being employed to evaluate it, and (b) that condition is opaque. These two characteristics together ensure that the project is procedurally self-defeating; and no such project can deliver a warrant for its outcome.

The two characteristics of Wright's revised template emphasized above—warrant-circularity and subjective indistinguishability—correspond (roughly) to these two conditions. (a) above just is an expression of warrant-circularity. (b)—opacity—is closely related to indistinguishability. But indistinguishability is both too weak and too strong. It is too weak because, as Wright describes it, indistinguishability requires only that the process/state of accomplishing the relevant putative warrant for P is subjectively *compatible* with R's holding (where R is itself inconsistent with a presupposition). This only requires, for example, that "it is disguised to look like a zebra" (R, the negation of the presupposition that it is not so disguised) *could* deliver the same experience to Bob as is delivered when an untampered-with zebra is in front of him. Opacity, by contrast, requires that the condition's failure *implies* that the same outcome will result. It's being a disguised zebra is not merely compatible with its looking to Bob as though it's a zebra; it *will* look like a zebra. That's what it means to say that it is (cleverly) so disguised. Consultation of those appearances will then deliver the outcome that it is a zebra once again.

Indistinguishability is also too strong, in that what's required for procedural self-defeat isn't that the opaque condition's failure generates *precisely* the same subjective experience—that, as Wright puts it, "things could be with one in *all* respects exactly as they subjectively are"—but only that they are similar enough, in *relevant* respects, so that implementation of the procedure will still deliver the same outcome. Zebras' stripes are not always (or ever) identically arranged; so if the zebra had instead been a disguised mule, the disguise would still count as such if the faux stripes were not arranged as they actually are. The condition that it not be so disguised—and so not disguised in that precise fashion—remains an opaque condition of warrant of the animal identification project.

So while the two characteristics Wright has emphasized point in the right direction, they—in particular, indistinguishability—have not quite arrived. Their true relevance, I suggest, is captured in the concept of a cognitive project's opaque enabling condition and the fact that, when the project's outcome enunciates such a condition, that project is procedurally self-defeating.

XVII. Procedural Self-Defeat and Dretske Cases

Now consider a variation of the station attendant case. You ask the attendant for directions downtown, and he asserts D as before. You now believe D. You then infer from D that M: it's not the case that downtown is some other way and the attendant asserted D order to mislead you.

Notwithstanding this difference, I suggest that the example still satisfies (a) and (b). The target of the project is M. This target is evaluated by a two-stage procedure: first, ask the attendant the way downtown, and second, infer whatever you can from his answer concerning M's truth-value.⁵⁹ D is the outcome of the first stage, which constitutes a sub-project of that overall project. M is a condition of that sub-project: if the attendant is misleading you, then you cannot learn the way downtown in virtue of his answer.⁶⁰

M is also a condition of the overall project generating M itself. In order to see this, note that it is a condition of any project designed to produce a warrant for its outcome by inference that the premises of the inference are themselves warranted.⁶¹ It is therefore a condition of the project involved in acquiring a warrant for M by inference from D that D be warranted. However, M is itself a condition of the project delivering a warrant for D. So M is a condition of the project delivering a warrant for D, and D's being warranted is a condition of the project delivering a warrant for M. M is therefore a condition of a condition of the project delivering a warrant for M, and so also a condition of its own warrant.⁶²

M is, moreover, an opaque condition of the sub-project involved in the first stage. M's being false—the way downtown is some other way entirely and the attendant asserted D in order to mislead you—implies that the attendant still asserted D. The sub-project's procedure—treating his assertion as

⁵⁹ This is, of course, a bizarre procedure for the evaluation of that target. We are in the process of explaining why it is bizarre.

⁶⁰ M's falsehood implies that the way downtown is some other way than it is in fact; that is why the agent, knowing the way downtown, would say D if M were false. So M's falsehood is incompatible with the facticity of your knowing D (by testimony). But that is, for the present point, irrelevant: if the attendant is attempting to mislead you, then you can't learn the way downtown no matter what he says, and no matter which is the way downtown.

⁶¹ Insofar as this is controversial, I bracket that controversy here.

⁶² More precisely, M is a condition of its own project on the assumption that D was itself produced by the project that in fact produced it (namely, asking the attendant). Some might balk at the "transitivity of project-conditions" being invoked. There are delicate issues concerning the individuation of projects that I cannot begin to explore here. It doesn't matter, however, for present purposes. Even if M is described more conservatively as an opaque condition-of-a-sub-project-whose-successful-outcome-realizes-a-condition-of-an-inferential-project-delivering-M, that opacity carries over in such a fashion as to ensure that the overall project is procedurally self-defeating: if M is false—and so the attendant is misleading you—he will still assert D and you will still utilize D as premise in the inference to M.

accurate—will still generate D as its outcome. Since the second stage involves deriving whatever can be derived from the result of the first stage concerning M's truth-value, the fact that D will be generated at the first stage when M is false implies also that M will be generated at the second stage, and so by the overall project. So the outcome of the project—M—is also an opaque condition of the project mobilized to deliver M. As before, this ensures that the project is procedurally self-defeating, and so cannot deliver a warrant for M.

This version of the testimony case is of course an analog for—indeed, an instance of—the Dretske cases we have been considering. All such projects involve inferences whose conclusions enunciate opaque conditions of the very projects employed to deliver warrant for those conclusions. Those projects are therefore procedurally self-defeating. If this is correct, then the problematic character of the inferences involved in these cases does not ultimately depend on the fact that they *are* cases in which inference is involved. They are a species of the genus of projects whose outcomes express one of their own opaque conditions, which is a genus of the family of procedurally self-defeating projects, which is a family of the order of self-defeating projects, all of which fail to deliver a warrant for their outcomes.

This remains entirely compatible with the success of the first-stage project involved in delivering a warrant to the premise in Dretske cases. D is not a condition of the project involved in asking the attendant the way downtown. Nor, therefore, is it an opaque such condition. So that premise may well be warranted (assuming, of course that the conditions for that warrant are in fact realized). But M, inferred from it, cannot be warranted in virtue of that inference. Although the premise is (true and) warranted, and the conclusion follows from it, the inference from those premises, warranted in the manner in which they are warranted, cannot deliver a warrant to that conclusion. So warrant transmission fails: the premise is warranted, but cannot transmit warrant to the conclusion. Call this the “procedural self-defeat explanation” (PSDE) of transmission failure.

XVIII. PSDE and Antecedent Warrant

PSDE has nothing to do with whether the agent is antecedently warranted in believing the conclusion. That C is a condition of project P with target T does not imply that the agent must possess an antecedent warrant for C in order for the project to succeed. Nor does C's being an opaque such condition imply this. Nor, therefore, does C's being an opaque condition of a project whose target concerns the realization of that very condition. The problem with these inferences is that they constitute cognitive projects whose procedures ensure that they will deliver the outcomes they do deliver

when those outcomes are false. It simply doesn't matter to that explanation whether the agent is already warranted in believing that outcome.

PSDE also does not imply that every warrant-circular argument fails to transmit warrant. More generally, it does not imply that any project whose outcome expresses a condition of that very project fails to deliver a warrant for that outcome. Such conditions might not be opaque. If not, they lack a crucial feature required by PSDE.

For example, it is arguably a condition of any project involving acquisition of empirical information that our senses are reliable (within certain respects and degrees, of course).⁶³ From some of that information, concerning the operation of our senses within the surrounding environment, we infer that our senses are indeed reliable (again within certain respects and degrees). That conclusion therefore expresses a condition of the project involved in delivering that very conclusion. However, it is not an opaque condition of that project. The falsehood of that condition—our senses are unreliable—certainly does not imply that the same outcome will result from implementation of that project. It *could* result; we could be brains in vats, in which case we would (appear to) proceed with the project and arrive again at the outcome that our senses are reliable. It may even result in the nearest possible worlds in which our senses are unreliable (although I see no reason to believe so). But the failure of the condition itself does not imply that it does so. PSDE does not therefore apply to this case.⁶⁴

This is appropriately analogous to the argument utilized by advocates of rule-circular arguments. As they point out, the availability of rule-circular inference in general does nothing, in itself, to ensure that such an argument in favor of any particular rule can be constructed. It is compatible with the permissibility of rule-circular arguments, for example, that our experience is sufficiently chaotic as to preclude the collection of an adequate sample of past apparently successful inductions from which to induce the reliability of induction.⁶⁵

By contrast, in the shabby zoo case the outcome (“it is not a disguised mule”) is not just a condition of the project involved in delivering its own warrant, it is an opaque such condition. It therefore constitutes a

⁶³ I remind the reader that its being such a condition implies nothing about the agent's awareness of its status as such.

⁶⁴ Analogous reflections apply to Wright's sheep-counting case.

⁶⁵ To be clear, I'm not taking a stand here on the success or otherwise of rule-circular arguments or the perceptual reliability argument (or on what that success might require). The point is only that it is compatible with PSDE that they do succeed. PSDE does not require that *only* inferential projects that are procedurally self-defeating fail to transmit, but that at *least* they do. In light of the fact that there may well be other forms of self-defeat realized by attempts at warrant-acquisition by inference, this is appropriate.

procedurally self-defeating project, which therefore fails to deliver a warrant for that conclusion. This result is independent of its being a shabby (vs. reputable) zoo, and so of whether Bob does or needs to possess an independent, antecedent warrant for its conclusion, as per the intuitions canvassed in sections VIII-XI above.

PSDE cannot, therefore, ground a defense of knowledge closure in the manner of the Wrightean strategy. Nothing in the applicability of PSDE to a particular case implies that the agent has, or will have, warrant for the conclusion by means other than transmission from the premise. It is therefore entirely consistent with a PSDE explanation of transmission failure in a particular case that the agent knows, and so is warranted in believing, the premise but is not warranted, and so does not know, the conclusion, and so that closure fails.

XIX. PSDE and Sensitivity

PSDE explains transmission failure by appeal, ultimately, to the intuition that no cognitive project can succeed when the procedure employed to evaluate the target itself ensures that the outcome delivered by the project will also be delivered if that outcome is false. It is hard to see how to defend that intuition further, except to re-emphasize the futility of such projects. At some point intuition bottoms out, at least without further appeal to a particular theoretical standpoint.

But some might worry that this explanation does rely upon such a standpoint and, in particular, upon one to which advocates of knowledge closure are already disinclined.⁶⁶ Sensitivity theories, very roughly, view the question what would happen, were the project's outcome false, as the epistemological linchpin: if that outcome would still be delivered if it were false, then the project fails to deliver a warrant for, and therefore knowledge of, that outcome.⁶⁷

Such a concern would, however, be misguided. To say that a project is procedurally self-defeating is not to say that, were the outcome false, the project would still produce that outcome. It is much stronger than that. Procedural self-defeat implies that it *follows* from the outcome's falsehood that

⁶⁶ They are so disinclined because classical versions of sensitivity theory imply closure failure. Not all sensitivity theorists do deny closure, however. Black 2008, for example, advocates a version of sensitivity consistent with closure. And DeRose's 1995 contextualist rendering of the sensitivity account is designed in part to avoid closure denial. (However, this essay's argument undermines the view that these versions of sensitivity are superior to the classical versions that do imply closure failure.)

⁶⁷ I'm purposely ignoring a multitude of variations on the sensitivity theme here. (See Becker & Black 2012 for some of them.)

that outcome will still result from that procedure.⁶⁸ If the agent is attempting to mislead when saying “I am not attempting to mislead”, then it follows that he will say this; he would not otherwise mislead. Assuming you trust him—doing so is your procedure—then you will believe he is not attempting to mislead. If the heliocentric theory is true—containing within its model our vantage point when looking up, as it must in order to explain the apparent celestial motions—then it follows that the apparent motions are as they are in fact. To infer from the apparent to the actual motions—and so to the conclusion that the sun rotates around the earth—will then imply that the heliocentric theory is false.

Similarly, if the animal is a mule disguised to look like a zebra, it follows that it looks like a zebra. If you use the two-stage procedure—determine what animal it is on the basis of its appearance and then infer from this whatever follows concerning the truth of “it is not a mule disguised to look like a zebra”—you will arrive again at the conclusion that it is not a mule disguised to look like a zebra. And similarly for the other Dretske cases.

Of course, if it follows from the falsehood of the outcome that the project will still generate that outcome, then it is also true that, if the outcome were false, the project would still generate that outcome. Detractors might claim that PSDE still borrows whatever plausibility it has from sensitivity as a result.

That would be a mistake. To see why, consider a (purported) case of insensitive knowledge acquisition, one of knowledge by inductive generalization.⁶⁹

Having noted that each among the multitude of past observed humans has, after a period of time, died, we induce that all humans are mortal. If, however there were an immortal human we would arguably still induce as we do. The falsifying immortal’s existence is, after all, compatible with our sample.⁷⁰ So in the nearest world—that most like the actual—compatible with the relevant antecedent we still have the same evidence, and so induce

⁶⁸ But in what sense of “follow”? In many cases the relevant sense is straightforward deductive implication: “it is disguised to look like a zebra” implies “it looks like a zebra” which, in conjunction with “it is what it looks like”, implies “it’s a zebra”. In other cases—the Brahe and attendant cases, for example—the inference is more elaborate, and may involve appeal to theoretical background and/or conceptual relations. There may as a result be subtle variations on this form of procedural self-defeat.

⁶⁹ The example is Wright’s, utilized in the development of the revised template; see Wright 2011, 92. But he is not alone in thinking that inductive generalizations create difficulties for sensitivity accounts. See, for example, Vogel 2007.

⁷⁰ He is a hermit, born in ancient Greece, who has managed to elude observation all this time.

as we do; the sensitivity conditional is false.⁷¹ And yet we can know, on the basis of that sample, that all humans are immortal.⁷²

Perhaps. However, it does not follow from the mere *existence* of an immortal that our investigation will not reveal him, and so that we will induce as we do. If this were so—if somehow the very procedure we utilize ensures that he and others like him will elude our observation, so that that procedure is guaranteed not to encounter a falsifying instance and we will induce as we do—then the intuition that such an induction delivers knowledge rapidly dissolves.

That is, however, what happens when a project exhibits procedural self-defeat: the outcome's being false implies that the project will still generate that same outcome. Its being a disguised mule implies that the procedure will deliver the result that it is not a disguised mule; any support for the proposition that it is a disguised mule must come from another source. Whether or not insensitivity is an epistemological sin, procedural self-defeat is a mortal one.⁷³

XX. Wright Back to Dretske

We have arrived at an explanation that might sound familiar to the cognoscenti of the literature on knowledge closure; it echoes a theme heard repeatedly in Fred Dretske's work on the issue. Here is a representative recent sample:

Though we rely on a measuring instrument to indicate the value of Q—that it is, say, 5 units—the measuring instrument does not indicate, does not carry information, that Q is not being misrepresented. The position of a pointer (pointing at “5”) on a well functioning meter indicates that the value of Q is 5 units, but it does not thereby indicate that the instrument is not broken, not malfunctioning, not misrepresenting a value of 2 units (say) as 5 units. Even when instruments (and this includes the human senses) are in perfect working order, they do not—they cannot—carry information that what they are providing is genuine information and not misinformation. That isn't an instrument's job. The job of a Q-meter is to tell us about Q, not about its reliability in providing information about Q.⁷⁴

⁷¹ So the argument goes. I myself am not convinced by it. But this is not the place to discuss the matter.

⁷² If you think we don't know this that way, substitute a different inductive generalization. The same argument applies.

⁷³ In an earlier version of this paper I myself attempted to explain transmission failure in subjunctive, sensitivity-like fashion. Thanks go to James Pryor, whose thoughtful comments on that version put me on the road to PSDE (which, I hasten to add, is not to say that he endorses it).

⁷⁴ Dretske 2005, 22.

The correlation between passages like this and Wright's recent comments is striking:

[Y]ou cannot, merely by looking at the animal, or the wall, get in position rationally to claim that conditions are suitable for the identification of animal species, or surface colours, just on the basis of unaided vision. That conditions are so suitable is a general authenticity-condition for each of the two projects concerned. But so are the more specific consequences of that condition detailed in their respective 3-propositions ["it is not a disguised mule", for example]. Neither the general condition, nor the specific forms of it detailed in the relevant 3-propositions, are open to merely visual confirmation.⁷⁵

Dretske offered his comments in way of denying closure; Wright offered his instead in way of denying transmission. Of course, even if attempts to learn about the project's presuppositions—about the reliability of the instrument—from the deliverances of the project itself—on the basis of the instrument's own readings—cannot succeed, that does not imply that we do not know those presuppositions anyway. It does not imply in particular that we don't need such knowledge to begin with, in order to learn anything from those projects and instruments.

On the other hand, it doesn't imply that we do. It is in fact somewhat more difficult to read the Wrightean strategy into Wright's later work on the topic.

For example:

Any cognitive project will be associated with a more-or-less open-ended set of authenticity conditions. And some of these... will be such that one can recognize in advance that the execution of the project will do nothing, except coincidentally, to rationally raise one's credence that they are satisfied. A proper subset of the latter may be included among the logical consequences of P [the premise arrived at by implementation of that project]. Should the project find in favor of P, the claim to warrant for P thereby afforded will not transmit to any among these... consequences of it."⁷⁶

The suggestion is that the recognizability-in-advance that the project can't raise one's credence in those presuppositions is the key to the explanation. There is at least an aura of PSDE here, since it is recognizable in advance that the project will deliver the same outcome when that outcome is false,

⁷⁵ Wright 2012, 468.

⁷⁶ Wright 2012, 470.

and so that the inference will be procedurally self-defeating.⁷⁷ At any rate, nothing in what Wright says here obviously appeals to the need for possession of advance warrant for those presuppositions.

So perhaps Wright is no longer an advocate of the Wrightean strategy, but advocates instead, or is at least moving toward, a PSDE-style explanation. If the argument of this paper is correct, he's moving in the right direction.

To do so, however, is to renounce the attempt to defend knowledge closure by appeal to the nature of transmission failure. As per section VIII above, PSDE—unlike the explanation offered by the Wrightean strategy—does not imply that closure succeeds whenever transmission fails. It therefore provides no comfort to the closure advocate that, when their intuition concerning cases in which transmission succeeds no longer applies, something else steps in to maintain closure. The oft-cited intuitive grounding for closure—in light of which it is treated by many as axiomatic in the theory of knowledge—falls away in precisely those cases that Dretske and others cite as counter-examples.

XXI. Conclusion

Condition (iv) of the Wrightean Strategy in defense of KC—that antecedent warrant for the conclusion is the only plausible explanation for transmission failure in Dretske cases—is false. Not only is it not an explanation at all, there is a plausible alternative explanation provided by PSDE.⁷⁸ The latter, moreover, does not imply KC when transmission fails. So the Wrightean strategy for the defense of KC fails.

It would be illegitimate to appeal at this point to KC itself in way of arguing that antecedent warrant for the conclusion is a condition of warrant for the premise in the relevant cases. If it is conceded that (i) WTP, encoding the transmission-intuition underlying the advocacy of closure, admits of exceptions and that (ii) the explanation for those exceptions, grounded in PSDE, does not preserve KC, then there is no longer any incentive to endorse KC in those cases. The rationale for endorsing KC itself as exceptionless, and so as applying in these cases as well, is lost.

⁷⁷ Nothing, however, in PSDE appeals to the fact that this is recognizable in advance. Note also that Wright speaks of claims to warrant rather than warrant itself, a shift I once again ignore for present purposes.

⁷⁸ Dretske cases also manifest a certain objectionable audacity: their conclusions declare that their own premises were respectably produced. In an earlier version of this paper I suggested that this was the key to transmission failure. And I still believe that it is why they manifest a presumptuous aura. But their audacity is only objectionable because they are in no position to make the claim they presume to make; that is, because they are procedurally self-defeating.

None of this demonstrates that closure fails. It is consistent with everything above that, whenever transmission fails, warrant for the premise does in fact require prior warrant for the conclusion (or that knowledge of the conclusion is assured for some other reason).⁷⁹ But the latter is not the explanation for the former. And PSDE, which does explain transmission failure, in no way ensures that the agent winds up knowing the conclusion. So those who concede that transmission fails in Dretske cases might as well deny closure too.⁸⁰

References

- Alspector-Kelly. (2011) “Why Safety does not Save Closure”, *Synthese* 183 (11): 127–142.
- Becker, Kelly & Black, Tim. (2012) *The Sensitivity Principle in Epistemology* (Cambridge: Cambridge UP).
- Black, Tim. (2008) ‘Defending a Sensitive Neo-Moorean Invariantism’, in V. F. Hendricks and D. Pritchard (eds.), *New Waves in Epistemology* (Basingstoke: Palgrave Macmillan), 8–27.
- Braithwaite, Richard. (1953) *Scientific Explanation* (Cambridge: Cambridge UP)
- Cohen, Stewart. (1988) ‘How to Be a Fallibilist’, *Philosophical Perspectives*, 2: 91–123.
- (2000) ‘Replies’, *Philosophical Issues*, 10: 132–139.
- Coliva, Annalisa. (2012) ‘Moore’s Proof, Liberals and Conservatives: Is there a (Wittgensteinian) Third Way?’ in A. Coliva (ed.), *Mind*,

⁷⁹ Safety accounts of knowledge are often advertised as enjoying the advantage over their sensitivity predecessors that they preserve closure. It is, in fact, far from clear that they do; see Alspector-Kelly 2011. But suppose they do. Safety theorists nevertheless don’t endorse WTP (they are neo-Moorean, but not Moorean). Nor do such accounts employ the Wrightean strategy. Bob knows that it’s a zebra, for example, solely because his experience ensures that he does not believe so in nearby worlds in which it isn’t, which in turn requires that there are in fact no nearby worlds in which it’s disguised (and so none in which it’s a disguised mule). But it does not require that, in addition, Bob be warranted in believing that it’s not a disguised mule. Safety theories preserve the letter of closure (perhaps), but they violate its spirit. Their supposed advantage in preserving closure (even if they do) is superficial.

⁸⁰ There are, of course, other reasons for denying closure. Many are objections, not to closure per se, but to particular theories implying its failure. This is particularly so with respect to sensitivity accounts, which are thought by many to imply, not just closure failures, but embarrassing ones. Some others apply to any closure-denying account, such as the “abominable conjunction” (DeRose (1995: 27–29)). There is no room to discuss these here. However, if the supposedly irresistible force of the intuition backing closure is, in the end, misapplied, then many of these other objections lose their sense of urgency: closure-denying views simply face a suite of objections as do all other epistemological views.

- Meaning and Knowledge: Themes from the Philosophy of Crispin Wright* (Oxford: OUP), 323–351.
- Davies, Martin. (2009) ‘Two Purposes of Arguing and Two Epistemic Projects’, in Ravenscroft I (ed.), *Mind, Ethics, and Conditionals: Themes from the Philosophy of Frank Jackson* (Oxford: OUP).
- DeRose, Keith. (1995) ‘Solving the Skeptical Problem’, *Philosophical Review*, 104: 1–52.
- Dretske, Fred. (1970) ‘Epistemic Operators’, *Journal of Philosophy*, 67: 1007–23.
- (2005) ‘The Case Against Closure’, *Contemporary Debates in Epistemology* (Oxford: Blackwell), 13–26.
- Hanson, N. R. (1958) *Patterns of Discovery* (Cambridge: Cambridge University Press).
- Hawthorne, John. (2004) *Knowledge and Lotteries* (Oxford: OUP).
- (2005) ‘The Case for Closure’, *Contemporary Debates in Epistemology* (Oxford: Blackwell), 26–42.
- Klein, Peter. (1995) ‘Skepticism and Closure: Why the Evil Genius Argument Fails’, *Philosophical Topics*, 23: 213–236.
- (2004) ‘Closure Matters: Epistemic Skepticism and Easy Knowledge’, *Philosophical Issues* 14: 165–184.
- Luper, Stephen. (2006) ‘Dretske on Knowledge Closure’, *Australasian Journal of Philosophy*, 84: 379–394.
- Nozick, Robert. (1981). *Philosophical Explanations* (Oxford: OUP).
- Plantinga, Alvin. (1993a) *Warrant: The Current Debate* (Oxford: OUP).
- (1993b) *Warrant and Proper Function* (Oxford: OUP).
- Pryor, James. (2004) ‘What’s Wrong with Moore’s Argument?’, *Philosophical Issues* 14: 349–78.
- (2012) ‘When Warrant Transmits’, in A. Coliva (ed.), *Mind, Meaning and Knowledge: Themes from the Philosophy of Crispin Wright* (Oxford: OUP).
- Pritchard, Duncan. (2007) ‘How to be a Neo-Moorean’, in S. Goldberg (ed.), *Internalism and Externalism in Semantics and Epistemology* (Oxford: OUP).
- Psillos, Stathis. (2004) ‘Inference to the Best Explanation and Bayesianism,’ in F. Stadler (ed.), *Induction and Deduction in the Sciences* (Dordrecht: Kluwer), 83–91.
- Silins, Nicholas. (2005) ‘Transmission Failure Failure’, *Philosophical Studies* 126: 71–102.
- Vogel, Jonathan. (2007) ‘Subjunctivitis’, *Philosophical Studies*, 134: 73–88.
- Williamson, Timothy. (2000). *Knowledge and its Limits* (Oxford: OUP).
- Wright, Crispin. (1985). ‘Facts and Certainty’, *Proceedings of the British Academy*, 71: 429–472.

- (2000) ‘Cogency and Question-Begging: Some reflections on McKinsey’s Paradox and Putnam’s proof’. *Philosophical Issues*, 10: 140–163.
- (2002) ‘(Anti-)Sceptics Simple and Subtle: G. E. Moore and John McDowell’, *Philosophy and Phenomenological Research*, 65: 330–348.
- (2004) ‘Warrant for Nothing (and Foundations for Free?)’, *Proceedings of the Aristotelian Society, Supplementary Volume*, 78: 167–212.
- (2007) ‘Perils of Dogmatism’. In S. Nuccetelli (Ed.), *Themes from G. E. Moore: New Essays in Epistemology and Ethics* (Oxford: OUP), 25–48.
- (2011) ‘McKinsey One More Time’, in A. Hatzimoysis (Ed.), *Self-Knowledge* (Oxford: OUP), 80–104.
- (2012) ‘Replies, Part IV: Warrant Transmission and Entitlement’, in A. Coliva (ed.), *Mind, Meaning and Knowledge: Themes from the Philosophy of Crispin Wright* (Oxford: OUP), 451–486.