

Some Reflections on the Acquisition of Warrant by Inference

Crispin Wright

I

A valid argument is one thing. A valid argument with warranted premises is a second. But a *cogent* argument is yet a third: it is an argument, roughly, whereby someone could/should be moved to rational conviction of the truth of its conclusion—a case where it is possible to *learn* of the truth of the conclusion by getting warrant for the premises and then reasoning to it by the steps involved in the argument in question. Thus a valid argument with warranted premises cannot be cogent if the route to warrant for its premises goes—of necessity, or under the particular constraints of a given epistemic context—via a prior warrant for its conclusion. Such arguments, as we like to say, ‘beg the question’.

Say that a particular warrant, *w*, *transmits* across a valid argument just in case the argument is cogent when *w* is the warrant for its premises. I do not know if anyone had registered the distinction between transmission of warrant, so characterized, and *closure* of warrant before I drew it in the 1985 British Academy lecture to which Martin Davies refers in his contribution to this volume. Closure of warrant across (known) entailment has of course been very widely discussed.¹ It is the weaker principle. A valid argument complies with closure provided that if there is warrant for its premises, there is warrant for its conclusion too. But a valid argument is transmissive if, roughly, to have warrant for its premises and then to recognize its validity is to acquire—perhaps for the first time—a warrant to accept the conclusion. My concern here will be with two forms of counterexample to transmission. Such cases need not be counterexamples to closure. Closure will hold but transmission may fail in question-begging cases—cases where there is warrant for the premises in the first place only because the conclusion is *antecedently*

warranted. (I am, myself, skeptical whether there are any genuine counterexamples to closure but that issue is not on our agenda.)

Note that transmission of warrant need not be an absolute characteristic of a given valid argument. There are at least two potential sources of relativity. First, it may be that the argument is such that one type of possible ground, w_1 , for its premises is transmissible—can yield a novel reason for accepting the conclusion when taken in conjunction with recognition of the validity of the inference—while another, w_2 , is not, but can only be possessed in the first place by a thinker whose information already includes warrant to accept the conclusion. Second, it may happen that, even when we focus on a specific warrant, w , the question whether that warrant transmits turns on the collateral epistemic context: that some but not all possible ways of acquiring w travel through the acquisition, or anyway prior possession, of warrant for the conclusion of the argument in question. Thus it may be that whether a given warrant transmits depends, as it were, on who wants to know—on who is to be persuaded of the conclusion and what their standards and presuppositions are.

The transmissibility of warrant is what makes for the possible advancement of knowledge, or warranted belief, by reasoning. It is only because warrants are sometimes (usually?) transmitted that deductions (other than those which discharge all their premises) are of use for anything other than the disclosure of commitments. But warrants are not always transmitted. When are they not?

II

A large and important class of nontransmissible warrants connect with the holism of empirical confirmation emphasized by Quine in the last two sections of “Two Dogmas.” This holism is pervasive. It involves that empirical confirmation is not a simple dyadic relationship but characteristically depends upon collateral information. At work at my desk in Philosophy Hall, I hear a thunderous rumble and sense a vibration in the building. Is that evidence of an incipient electric storm? Yes, if the sky has darkened and the atmosphere is heavy and still. Probably not, if the sky outside is clear blue and my office overlooks Amsterdam Avenue with its regular cargo of outsize trucks. I see a substantial-looking brownish bird of prey perching on a fence post. A sighting of a Golden Eagle, perhaps? Quite possibly, if I am in the north-western Scottish Highlands; but not if I am knowingly in Welsh farmland, where buzzards

are now quite common. Examples such as these suggest that what is normal in empirical cases is *information dependence* of warrant. A body of evidence, *e*, is an information-dependent warrant for a particular proposition *p* if whether *e* is correctly regarded as warranting *p* depends on what one has by way of collateral information, *I*. Consider a case where one's collateral information, *I*, does indeed sustain *e*'s warranting *p* but where *e* could not rationally be regarded as warranting *p* if certain elements of *I* were missing and uncompensated for. Such a relationship is always liable to generate examples of transmission failure: it will do so just when the particular *e*, *p*, and *I* have the feature that needed elements of the relevant *I* are themselves entailed by *p* (together perhaps with other warranted premises.) In that case, any warrant supplied by *e* for *p* will not be transmissible to those elements of *I*. Warrant is transmissible in such a case only if a rational thinker could cite as her ground for accepting *I* the fact that she has warrant for *p*, supplied by *e*, together with the entailment. No rational thinker could do that if the warrant for *p* supplied by *e* originally depends on prior and independent warrant for *I*.

It is easy to generate examples of transmission failures under this general template (*the information-dependence template*). For instance, (AIRPORT):

You are waiting in an airport lounge and,

(*e*) You hear the agent utter the words, "This is a final boarding call for Northwest's flight NW644 to Minneapolis."

So you naturally infer that,

(*P*) The agent has just orally forewarned passengers in English of final boarding for NW644.

P entails *I*:

(*I*) The agent understands (some of) a language (English).

But clearly the warrant bestowed on *P* by *e* does not transmit across this entailment from *P* to *I*. Rather it is only in a context of collateral information in which *I* is justifiably assumed that *e* provides a warrant for *P* in the first place.

Or consider (TWINS): Jessica and Jocelyn are identical twins whom you know well but have difficulty distinguishing.

Suppose,

(*e*) You see a girl approaching you who looks just like Jessica.

There is a defeasible inference from that to,

(*P*) That girl is Jessica.

And an entailment from there to,

(*I*) That girl is not Jocelyn.

But given your discriminatory limitations, there is no question of treating *e* as a warrant for *P* and then transmitting it across the entailment to conclude *I*. Rather you—though not perhaps someone who can distinguish the twins purely visually—will need the latter already in place as collateral information before you can reasonably take *e* as a warrant for *P*.

III

As formulated, the information-dependence template engages only the transmission of *inferential* warrants: warrants consisting in the possession of evidence, *e*, which licenses a defeasible inference to one or more of the premises of the entailment in question. But what of noninferential cases—cases where warrant for a premise is acquired not on the basis of evidence but directly, via the operation of some cognitive faculty—perception, or memory (on some construals), or logical or mathematical intuition, or perhaps a faculty of immediate self-knowledge—which we regard as directly responsive to the subject matter in question? Can such, noninferential warrants fail to transmit? If so, when and why?

Here are two famous examples in modern epistemology.² You go to the zoo, see several zebras in a pen, and opine (ZEBRA): that those animals are zebras. Well, you know what zebras look like, and these animals look just like that. Surely you are fully warranted in your belief. But if the animals are zebras, then it follows that they are not mules painstakingly and skillfully disguised as zebras. Does your warrant transmit to the latter claim? There is a strong intuition that it does not. Did you examine the animals closely enough to detect such a fraud? Almost certainly not. The grounds you have for (ZEBRA)—essentially, just the look of the beasts—have no bearing on this possibility.

Again: you look at a wall and see that it is painted red. So you have acquired a warrant for thinking (RED): that it is red. But its being red entails that it is not a white wall cleverly illuminated by concealed light-

ing to look as if it is red. So have you thereby acquired a warrant for thinking *that*? Again, the strong intuition is not. Your warrant was acquired just by looking at the wall—no doubt you did enough to verify that it is red if indeed it is, but what you did simply didn't reckon with the possibility of deceptive concealed illumination.

When, in the contexts described, you form your beliefs about the zebras and the color of the wall, there are external preconditions for the effectiveness of your method—casual observation—whose satisfaction you will very likely have done nothing special to ensure. Made-up mules and tricky lighting involve the frustration of those preconditions. Can the warrants you acquire licitly be transmitted to the claim that those preconditions *are* met,—or at least that they are not frustrated in those specific respects? It should seem obvious that they cannot. While you have—no doubt quite justifiably—taken it for granted that the conditions were generally suitable for the acquisition of reliable information by casual-perceptual means, it would be absurd to pretend that you had *gained a reason for thinking so*—at least in the specific respects that you didn't have to reckon with disguised mules or deceptive lighting—just by dint of the fact that those specific possibilities are logically excluded by the beliefs which, courtesy of your background assumption, you have now confirmed.

Dretske, of course, originally presented these cases as failures of *closure*. I have just presented them as failures of transmission. But once the distinction is on the table—as it was not in Dretske's discussion—I think it is clear that the latter is the correct diagnosis. If they were cases of failure of closure, then it should be possible clear-headedly to claim a perceptual warrant for (ZEBRA) while simultaneously disclaiming all warrant for the proposition that the animals in question are not disguised mules. That would be to concede that it could be—for all one was entitled to suppose—that conditions are unsuitable for basing beliefs such as (ZEBRA) on casual perception, because it could be that the zookeeper has been cutting corners and disguising his animals, etc. *Mutatis mutandis* for (RED) and the red wall. So one ends up making a claim of the form: I have a particular kind of warrant for believing *P* but, for all I am entitled to suppose, it may be that conditions are unsuitable for getting a warrant for *P* of that kind. And that cannot be a clear-headed claim. If I assent to it, what can I possibly imagine entitles me to claim the warrant for *P* in the first place?

However, if this diagnosis is right, then we do have to reckon with a range of examples of transmission failure in which information

dependence of warrant and the *e*, *P*, *I* structure is not—or at least not obviously—involved. At least at a first pass, the impression that there is a transmission failure in, e.g., the zebra case seems intuitive even when one takes it that perception is a direct engagement with the local environment and the warrants conferred by it are, in the basic case, not inferential but direct.

What form of diagnostic template do these examples suggest? Here is a suggestion I have made in other work.³ Suppose I take myself to have a noninferential warrant for a proposition *A*. And let the question be whether this warrant transmits to a certain consequence, *B*, of *A*. However, suppose *A* stands opposed to some proposition *C* whose truth would undermine my warrant for *A* and which could be true in certain circumstances *subjectively indistinguishable* from those in which I actually find myself. And suppose *C* does not entail *B*, but would be true if *B* were false.

The set-up is thus (THE DISJUNCTIVE TEMPLATE):

- (i) that *A* entails *B*;
 - (ii) that my warrant for *A* consists in my being in a cognitive state—perceiving, remembering, mathematically intuiting, or whatever—which is subjectively indistinguishable from a state in which the relevant *C* would be true;
 - (iii) that *C* is incompatible with the reliable operation of the cognitive capacities involved in generating the warrant for *A*;⁴
- and
- (iv) that *C* would be true if *B* were false.

Suppose I know all this. The key question is what, in the circumstances, can justify me in accepting *A*? Should I not just reserve judgment and stay with the more tentative disjunction, either (I have warrant for) *A* or *C*? For it is all the same which alternative is true as far as what is subjectively apparent to me is concerned. The answer has to be, it would seem, that the more tentative claim would indeed be appropriate unless I am somehow *additionally* entitled to discount alternative *C*. It may be that I have collateral information telling against *C*. Or it may be that, for one reason or another, I am not required to bother about *C*. But notice that either way, in order for me to be entitled to discount *C*, and so move past the disjunction to *A*, I have to be entitled to *discount the negation of B*, and therefore entitled to accept *B*; for by hypothesis, if not-*B*

were true, so would *C* be.⁵ So it would seem that I must have an appreciable entitlement to affirm *B already*, independent of the recognition of its entailment by *A*, if I am to claim to be warranted in accepting *A* in the first place. The inference from *A* to *B* is thus not at the service of addressing an antecedent agnosticism about *B*. So my warrant does not transmit.

It is straightforward to cast the zoo and red wall examples into this template. Let *A* be the proposition that the animals in question are zebras; *B* is accordingly the entailed proposition that they are not mules painstakingly and skillfully disguised as zebras; while for a suitable *C* we need look no further than a generalization of the negation of *B*: say, that the animals in question are not zebras but just look that way. *C*, so selected, meets condition (iv). It also meets condition (iii): clearly, if we are having to deal with circumstances in which animals' appearances are deceptive, then conditions are unsuitable for the reliable operation of the relevant cognitive capacities—those involved in the identification of animal species by casual observation of their appearance. Or again, let *A* be the proposition that the wall is red, *B* the entailed proposition that the wall is not a white wall cleverly illuminated by concealed lighting to look as if it is red, and *C* the proposition that it is not a red wall but looks just like one. *C* meets condition (iv). It also meets condition (iii): again, if we are having to deal with circumstances in which things' apparent colors are deceptive, then conditions are unsuitable for the reliable operation of the relevant cognitive capacities—those involved in the identification of color by casual observation of an item's appearance. However—condition (ii)—the experiences of seeing zebras and seeing mere zebra look-alikes are relevantly subjectively indistinguishable; as are the experiences of seeing a red wall and seeing a mere red-looking wall. So in treating my state as being a bona fide perception of zebras or a red wall respectively, I implicitly discount the uncongenial, deceptive alternatives *C*. And now, whatever my warrant for doing so, it has to be there *already*; and it must provide warrant for the respective instances of *B* independently of any consideration of their entailment by the corresponding choices for *A*.

IV

Now to content externalism and the McKinsey argument. In another essay,⁶ I have suggested that the disjunctive template can be applied to corroborate the idea, independently proposed by Martin Davies in the

course of recent work,⁷ that McKinsey's paradox—the apparent incompatibility of content externalism with our ordinary conception of self-knowledge—turns on the presumed transmission of what is in fact an untransmissible warrant. The thought, roughly, is that in the context of my acceptance a priori of a content externalism strong enough to sustain the type of proposition typified by McKinsey's second premise—which we can take as

McKinsey (ii) If I believe that water is wet (and hence have the concept of water), I belong to a speech community that has encountered water

—the holding of the conclusion:

McKinsey (iii) I belong to a speech community that has encountered water,

becomes something for which I presuppose a prior entitlement in taking it that I may justifiably claim to know the other premise:

McKinsey (i) I believe that water is wet,

in the fashion involved in normal noninferential self-knowledge. Why might one think this?

A content externalism strong enough to sustain McKinsey's second premise a priori requires not just that the *identity* of certain concepts is externally individuated—that my concept of water, for instance, varies under counterfactual variation in the actual identity of the watery stuff with which I and my speech community interact—but that the very *existence* of my concept of water turns on the existence of an appropriate extension for it: in short, that on Paul Boghossian's Dry Earth, for example⁸—in which, incredibly, all apparent interaction with watery substance is multisensory communal hallucination—I would simply have no water concept, even though my experience was subjectively indistinguishable from what it actually is. We are, in other words, implicitly entertaining types of content-externalism about common nouns which bear a relevant analogy to the strong referentialism about certain kinds of singular terms espoused by Russell and Gareth Evans. According to strong referentialism, referentless singular terms lack sense as well: they can play no role in fixing the truth conditions of a thought. On such a view, if I hallucinate a dazzlingly colored bird and say to myself, "I wonder if *that* is a kind of Oriole," I actually fail to express any thought thereby, even though subjectively my situation is just as if I were think-

ing a demonstrative thought. In like cases, according to the relevant kind of externalism, a thinker who lives (always or for long enough) on Dry Earth and who says to herself, "Water is wet," suffers a similar *illusion of content*, even though in a condition subjectively indistinguishable from her *doppelgänger* on Earth who thereby thinks the routine thought that water is wet.

It is this implicit provision for the possibility of content illusion which is crucial to the misgiving about transmission in McKinsey cases. Suppose that I (an English speaker) want to describe myself in the words of McKinsey's first premise: "I believe that water is wet." If some form of content externalism is true which is strong enough to sustain the second premise, then there are external preconditions of my expressing a true belief by those words—precisely the appropriate history of interaction with watery stuff in the world—whose satisfaction I may nevertheless, without compromise of the warrant for my claim, have done nothing special to ensure (just as I did—and normally would be required to do—nothing special to ensure the appropriately nondeceptive character of the situation in which I observed Dretske's purported zebras.) Does my routine psychological warrant transmit to the claim that those preconditions *are* met?

The comparison suggests not. But how do matters pan out under the disjunctive template? Take A as the proposition that I believe that water is wet, B as the proposition that I, or my speech community, has had such and such encounters with water, and C as the proposition that my tokening of "I believe that water is wet" is content-defective owing the reference failure of the purported natural kind term, "water," in my language. Then, *prima facie*, each of the four conditions required by the template is met:

- (i) Proposition A, that I believe that water is wet, entails—on the assumption of the conceptually necessary truth of the relevant strong externalism—proposition B, that I, or my speech community, has had such and such encounters with water.
- (ii) My warrant for A consists in my being in a state which is subjectively indistinguishable from a state in which the relevant proposition C, that my tokening of "I believe that water is wet" is content-defective owing the reference failure of the purported natural kind term, "water," in my language, would be true.
- (iii) C is incompatible with my having warrant for A in this world; for if C is true, there is no such thing (in this world) as the belief that

water is wet; and if I have warrant for *A*, then that proposition exists and is exactly what (in this world) “I believe that water is wet” expresses, so that sentence is not content-defective, contrary to *C*.

Finally,

(iv) *C* would be true if *B* were false—for if our encounters with water had never occurred, that would suffice, in the presence of the relevant strong externalism, to divest us of the concept water and thus to ensure that all purported expressions of it are content-defective.

The consequential analysis of the transmission failure involved in the McKinsey argument—always assuming cogent a priori motivation for its externalist premise—is thus simply that once that premise is known and in play, my routine psychological warrant for taking it that *A* is true—and hence for dismissing the uncongenial interpretation of my subjective state as one in which *C* holds—must depend on antecedent warrant to think that my tokenings of ‘water’ comply with appropriate externalist constraints, in particular as described by *B*. So in taking it that I am warranted in accepting *A*, I presuppose a warrant for *B* independent of the recognition of its entailment by *A*. The recognition of that entailment is thus not at the service of my learning *B*, and the paradox—that my (broadly construed) a priori knowledge of *A* and of the correctness of strong content externalism would lead to similarly but incongruously a priori knowledge of *B*—is dissolved.

One vital clarification is required. If the demanded antecedent reason for discounting *C* had to be *empirically* acquired, this line of thought would indeed involve abandoning the groundlessness (a priority) of self-knowledge in the relevant kind of case; rather, empirical evidence would be presupposed that ‘water’ indeed possessed the appropriate historical connections and McKinsey’s own original incompatibilist conclusion would then be sustained. Even then, if the diagnosis is right, there would be a transmission failure. But my contrary suggestion is this: that even if one effect of a correct externalism is indeed to introduce, via the possibility of illusions of content, a novel form of possible *defeat* for self-knowledge claims, the epistemological impact of that possibility ought to be conceived much as that of the possibility of perceptual illusion vis-à-vis the justification of perceptual claims. I propose, that is, that we may reasonably allow an *a priori* *presumption* against that possibility, since to allow an *empirical* issue to arise in every case would be to abrogate the means to resolve any such issues—I have to take certain perceptions to be reliable if I am to investigate the reliability of others;

I have to take certain seeming-thoughts to be well-founded if I am to investigate the well-foundedness of others. If that is broadly the right way to look at the matter, at least as a starting stance, then the groundlessness (a priority) of basic self-knowledge—properly conceived—is unimpugned by externalism. We have a (defeasible) entitlement to set aside the uncongenial *C* without evidence, and the effect of the McKinsey deduction is, not to make available a nonempirical warrant for its conclusion but to bring out the empirically unearned but justified presumption of its conclusion, and of the satisfaction of external conditions on content in general, on which first-person authority for the contents of ones attitudes—at last within a (sufficiently strong) externalist framework—must rest.⁹

These reflections may stand some elaboration, so let me draw out their bearing upon an objection lodged by Michael McKinsey in his contribution to the present volume. McKinsey very reasonably notes that the observation of transmission failure is not enough to draw the sting of his argument—for the paradox most basically relies not on the transmissibility of the combined warrant for its premises but on the *closure* of a priori warrant across (known) entailment. Let it be conceded that I cannot *learn* of my speech community's history of interaction with water merely by reasoning from my belief that water is wet and an appropriate form of content-externalism. Still the paradox will linger so long as we are still forced to acknowledge—by an appropriate closure principle—that it is possible to come to know from the armchair, as it were, that my speech community has interacted with water, even if this possibility cannot be realized by transmitting the warrant for the McKinsey premises across the reasoning of his argument but has to be accomplished in some other—unspecified—way.

Now, we already observed (note 9) that transmission failures actually presuppose closure—that if the warrant for a set of premises fails to transmit to a consequence of them, that will be because, broadly, an antecedent warrant for the consequence needs to be in place in order for a thinker to access the warrant in question for the premises, and hence that if (that) warrant for the premises exists, so will (some kind of) warrant for the conclusion. So closure of warrant per se is not in question in the cases that interest us. But McKinsey's point needs more than that. The crucial question is: If the premises of a valid argument are warranted a priori, of what kind of warrant for its conclusion is the availability thereby ensured? In particular, is the potentially troublesome closure principle, that closure holds for a priori warrant, sound?

McKinsey distinguishes strong—empirically indefeasible—and weak—empirically defeasible—notions of a priori warrant and argues in some detail that closure holds for both. I think it is clear that he is right. Closure holds trivially where transmission does. But what characterizes cases where transmission fails is that warrant for the conclusion is presupposed by warrant for the premises—that achieving warrant for the premises has to proceed through the achievement of warrant for the conclusion. So if the latter had to go via an empirical route, so would the former; conversely, any kind of a priori warrant for the premises would have to involve no less—so to speak—of an a priori warrant for the conclusion. So isn't that enough to reinstate the paradox, failure of transmission notwithstanding?

No, it is not. We need a further distinction. What would be paradoxical would be the idea that I could *earn* a warrant for—win through to knowledge of—the proposition that, e.g., I belong to a speech community that has encountered water purely by exercise of the broadly reflective means available to me in Davies's armchair. A transmitted warrant would be such an earned warrant, but McKinsey is quite right that closure for reflectively earned armchair warrant would suffice to set the problem up. However we don't have *that* principle—quite. What we have (according to my suggestion above) is merely an *a priori presumption* in favor of the integrity of the concepts in terms of which I essay to formulate items of my self-knowledge, which in conjunction with the known conceptual necessity of an appropriate kind of externalism becomes an a priori presumption in favor of the satisfaction of the external conditions necessary for that integrity. Such an a priori presumption—or *entitlement* (now to propose a specialized use of this term)—may be counted as a subspecies of warrant. But it is conferred not by positive evidence for the proposition in question but by the operational necessity, so to speak, of proceeding on the basis of such so far untested assumptions if one is to proceed at all. I have a similar entitlement, *ceteris paribus*, to assume the proper functioning of my perceptual apparatus on a particular occasion: To be sure, the matter could be empirically investigated but only in a context in which my own or other's perceptual apparatus was again assumed, untested, to be functioning adequately.

Notice, as the latter example brings out, that an a priori entitlement to a belief is quite consistent with the only envisageable kind of positive evidence for it being empirical. I can be a priori entitled to suppose that my senses are functioning adequately right now, though a

check would need to be empirical. Likewise—the proposal is—I can be a priori entitled to suppose that the seeming-concepts configured in a thought I seem to be having are all well-founded though—for externalism—their nature may be such that checking the point would demand some social and environmental history. But this is enough to save the reflective phenomenology of self-knowledge consistently both with externalism and with closure of a priori warrant, provided warrant may cover both earnings and entitlements. Once I accept the conceptual necessity of the appropriate externalism, my warrant for the thought, ‘I believe that water is wet’, does indeed presuppose—is accessed via—a warrant for my speech community’s historical interaction with water. But the latter warrant does not need to be earned by investigation. If it did, it would of course have to be earned empirically, and the warrant for my second-order belief would have to be acquired empirically as well. Entitlements, however, are available in the armchair. After externalism, self-knowledge of the contents of one’s attitudes is indeed opened to a new kind of empirical defeat. But it can remain achievable in the armchair manner allowed for by traditional ordinary psychology whenever there is an *a priori presumption that the relevant defeaters do not obtain*.

There is of course much more to say on the topic of entitlement. It is a topic of huge importance. But I hope I have said enough to make it seem plausible that the existence of a presumption in favor of the good standing of one’s concepts should be an entitlement if anything is.

V

The disjunctive template discloses that the McKinsey argument does not involve a transmissible warrant. And the impression that closure of a priori warrant is anyway enough for its incompatibilist lesson trades on overlooking the possibility of a priori entitlements for beliefs for which an earned warrant would have to be empirical. So goes the proposed dissolution of the paradox.

But there is a misgiving which needs to be confronted. It concerns the fourth condition of the disjunctive template—specifically the claim that when strong externalism is assumed, freedom by me and my speech community from all historical contact with water would suffice to induce content defectiveness into our purported water thoughts. That is certainly so when the alternative scenarios considered are of the Dry Earth type. But what justifies us—theorists—in restricting attention to those?

If Twin Earth scenarios are considered as well, then our purported water thoughts would suffer not divestment of content but *change* of content. Had our actual encounters been with twater, not water, the relevant *B* would have been false. But the relevant *C* would not have been true: my term, 'water', would not have been divested of content but would have expressed the concept, *twater*, instead. So the template, it seems, does not cleanly fit the case and the diagnosis of transmission failure is compromised.

What should we make of this? It is possible, of course, that the problem is in the detail of the template—or even that it is misconceived in some more radical way. But I suspect that the truth may be a little more interesting: namely, that there is a transmission failure in the vicinity, so to speak, of the McKinsey argument but that it is not exactly as proposed; and that when we see what it is, we will be able to confirm the spirit of the original diagnosis and dissolution of the associated paradox while granting that, strictly, the conjoined warrant for the premises in the McKinsey argument *does* transmit to the original conclusion after all. Let me explain.

Put aside the detail of the disjunctive template and reflect for a minute on the intuitive diagnostic thought which motivates this response to the McKinsey paradox. Strong externalism opens up the (doubtless merely theoretical) possibility of previously unenvisioned kinds of defeaters for claims about one's own mental states, connected with illusions of content. But the situation compromises our right to such claims in the kind of circumstances in which we normally make them only if we now need to do work—specifically, to do *empirical* work—to ensure those defeaters do not obtain in any particular case. My suggestion has been that in general we do not: that, absent reason to suppose the contrary, we may take it that we have a standing (though defeasible) entitlement to the suppositions that all is in order with our concepts, as it is with our vision and other cognitive faculties, that others' testimony is sincere, and that the appearances of objects around us are not systematically misleading. In general, all noninferential warrants are conferred subject to such background entitlements. So one important general limitation on transmission will apply wherever the conclusion of an inference spells out a known constitutively necessary condition for the realization of an entitlement which conditions the acquisition of a particular kind of noninferential warrant for one of its premises. If it can be agreed that the integrity of the relevant concepts is such a background entitlement for claims of all kinds, including claims concerning one's

own mental states, then it follows one cannot transmit warrants for such claims across entailments to conclusions that specify what we know to be constitutively necessary conditions for the integrity of concepts involved in those very claims.

If this basic idea is granted, it is inescapable *either* that there must be a transmission failure in the McKinsey argument, somehow not quite brought into focus by the disjunctive template, *or*—contrary to our assumption throughout—that the conclusion of the argument is not, after all, a necessary condition for the freedom from content illusion tacitly assumed as a background entitlement when I take it that I know McKinsey (i). Well, the awkwardness noted a moment ago about the fourth condition of the template now emerges, in effect, as a signal that it is the second alternative which holds: lack of water encounters would indeed *not* suffice for content illusion in my claimed item of self-knowledge (though it would impinge on the *identity* of what I claim). But if the basic diagnostic thought just adumbrated is correct, there will still be a transmission failure in the inference from that claim—advanced as a routine piece of noninferential self-knowledge—to *whatever* is a constitutively necessary condition for freedom from content illusion. So the question is: What is that necessary condition, as far as strong externalism and the (apparent) involvement of the concept of water is concerned?

Precisely, that one not inhabit Dry Earth. It is if everything which we would take to have been an actual encounter with a watery substance—water or an epistemic counterpart of water—is supposed to have been illusory that strong externalism will convict my purported claim to believe that water is wet of content illusion. So consider what happens if we run the McKinsey argument on one step further, so that it becomes (THE EXTENDED MCKINSEY ARGUMENT):

McKinsey (i) I believe that water is wet.

McKinsey (ii) If I believe that water is wet (and hence have the concept of water), I belong to a speech community that has encountered water.

McKinsey (iii) Therefore, I belong to a speech community that has encountered water.

McKinsey (iv) Therefore, it is not the case that I belong to a speech community that has never encountered either water or any other watery substance.

Now, both the disjunctive template and the more general diagnostic thought which motivated it will coincide in diagnosing a transmission failure in the inference from McKinsey (i) and (ii), via McKinsey (iii), to McKinsey (iv). If McKinsey (iv) were false, my psychological claim McKinsey (i) would indeed fall prey to content illusion. But it is McKinsey (iv) that constitutes the large empirical claim about the character of the world and our history of activity within it for which it would be paradoxical to claim that warrant can be earned just by ordinary reflection on one's beliefs and an a priori warranted content externalism. It is McKinsey (iv) that is "not the kind of thing that can be known in that way." To be sure, McKinsey (iii) *looks* just as adventurous, taken in isolation. But once McKinsey (iv) is taken to be an a priori entitlement, necessary to make good the local instance of the general assumption of conceptual integrity presupposed in all (self-)knowledge, the transmission of a nonempirical warrant as far as McKinsey (iii) adds to McKinsey (iv) only in respects which can indeed quite properly be viewed as nonempirical. For what McKinsey (iii) adds to McKinsey (iv) is only the identification of the watery stuff of our presumed acquaintance as *water*. And that *is* plausibly a priori. It is a priori that water is the watery stuff of our actual acquaintance, if there is any watery stuff of our actual acquaintance—compare: it is a priori that we are situated here, rather than over there, if we are situated anywhere at all.

VI

If I am right in this new diagnosis, then there is strictly no transmission failure in McKinsey's original argument after all. Relative to the entitlement to assume the integrity of the concepts drawn upon by any putative item of self-knowledge, the warrant for the claim that one believes that water is wet will—in conjunction with an a priori justification for strong externalism (if any exists)—transmit across the inference to the claim that one's speech community has a history of interaction with water. What it will not transmit to is the claim that one's speech community has a history of interaction with some watery substance. Rather, in an informational setting in which one's standing entitlement to assume the integrity of one's concepts—in particular, the concept *water*—already encompasses the latter claim, the nonempirical warrant transmitted across the McKinsey inference is merely for the identification of that substance as water. And that, I have suggested, is harmless.

In his recent work on these matters Martin Davies has been developing his own characterization of two kinds of limitation on warrant transmission. One of his principles effectively converges in its verdicts on those of the information-dependence template, and so far as I can tell, may well be just an alternative formulation of it. But the other—the Second Limitation Principle—stands somewhat in contrast to the proposal incorporated in the disjunctive template. This principle has undergone various reformulations in Davies's successive essays on the topic, but it has rejected transmission in the McKinsey argument in each of its incarnations. If my present diagnosis is right, and there is actually no such failure, then there must be some error in Davies's principle. Alternatively, if Davies is right, then the disjunctive template—and the background thought which drives it—must be too weak somehow, failing to detect the transmission failure that the McKinsey argument actually involves. Which is the fact?

Well, let me conclude on a note of tentative reservation about Davies's principle. His latest formulation (this volume) runs as follows:

Epistemic warrant cannot be transmitted from the premises of a valid argument to its conclusion if, for one of the premises, acceptance (i) of the assumption that there is such a proposition for the knower to think as that premiss, and (ii) of the warrants for the other premises, cannot be rationally combined with doubt about the truth of the conclusion.

Both in the essay from which this formulation is taken and in earlier work, Davies has shown himself very sensitive to the issue of motivation for his principle. I have to confess to feeling that there still is an issue about that. The principle requires that, if I reason from what I regard as warranted premises to a conclusion a doubt about which would be inconsistent with the availability of one of the premises to my thought, then I cannot thereby learn of the truth of that conclusion. But why not? Suppose I am troubled by some skeptical scenario—by my inability to exclude the apparent possibility that it might be true. If Davies is right, I cannot rationally come to realize that it is not true by providing warrant for the premises (i) that I grasp a proposition describing it, and (ii), that no creature in such a scenario could have access to the conceptual wherewithal to grasp that proposition, and then proceeding to reason in the obvious way from them. But intuitively there is nothing wrong with the *strategy* of such an argument—whatever one thinks of the detail of Putnam's famous implementation of it. So the Second Limitation Principle impresses as too strong.

Naturally this point will not impress anybody who is already convinced that the cogency of Putnam's Proof stands or falls with that of the McKinsey argument in any case. But there is a second, less tendentious misgiving about the strength of Davies's principle. It concerns its apparent insensitivity to something we registered at the start—the potential warrant relativity of issues having to do with transmission. Davies's Principle speaks of the rational uncombinability of three attitudes: doubt about the conclusion of an argument, acceptance of warrants for all (but one) of the premises, and acceptance of the availability of one in particular of the premises (the other one) to one's thought. But what kind of consideration is envisaged as determining that these three attitudes cannot rationally be combined? Well, what determines that in the McKinsey case is nothing, it seems, but the relevant attitudes' respective *contents*: I cannot rationally believe simultaneously that I have a belief which draws on the concept 'water', and that there is (all things considered) warrant to accept that anyone who has such a belief meets a certain historical condition while doubting—believing that it is not the case that—I meet that condition. I cannot do so simply because that is to doubt an *obvious consequence* of things I consider that I warrantedly accept. But if that is how the Second Limitation Principle is supposed to work, then it is going to block transmission in McKinsey's reasoning *whatever* the nature of the warrant accepted for its conditional premise—McKinsey (ii). And that is clearly a bridge too far. For McKinsey's argument is fine as a vehicle for warrant transmission provided the warrant for its major premise is *empirical*. Suppose, for instance, I am a dyed-in-the-wool antiexternalist who—perhaps as a result of a brain injury—lacks all recollection of (testimony of) encounters by myself or anyone in my speech community with water. Interested in what are the actual *empirical* preconditions for possession of the concept *water*, I do some interactive anthropology. An electronic interrogation of members of a variety of societies—not extending to my own—who possess the concept discloses that they all have a history of interaction with the stuff. I therefore propose on inductive grounds that this probably goes without exception, and so reason my way to the prediction that this will also prove true of my own speech community. *This*—*empirical*—warrant surely transmits. Thus McKinsey's argument is not a case of transmission failure *tout court* but at best illustrates the relativity we noted at the outset. Any diagnosis of transmission failure within the argument must therefore relate specifically to the setting when its second premise is warranted in the special fashion purportedly

provided for by externalism, and known to be so. This the Second Limitation Principle, at least as naturally understood, would seem to fail to do.

More generally, no restriction on warrant transmission can be generally satisfactory which is sensitive merely to rational cotenability relationships among attitudes determined *purely by their contents*. It remains to enquire, therefore, whether the disjunctive template fares better on this point. Can we use it to explain why reasoning—the very same reasoning—from the premises of the extended McKinsey argument to its conclusion may transmit warrant when the second, conditional premise is conceived as an empirical contingency but not when it is conceived as a conceptual necessity, holding a priori on externalist grounds?

One immediate difference, of course, is that in the former case the conditional premise—McKinsey (ii)—becomes essential to the deduction and we can no longer take McKinsey (i) and (iii) respectively as propositions A and B for the purposes of the template since, so identified, A does not entail B. But without here venturing to consider how best to generalize the template to cope with the multipremise case—work that would have to be taken on in a properly detailed treatment—we can still appreciate what I think is the crucial contrast. To recap. In no case can I rationally claim warrants for the premises of an argument unless I am entitled to take it that all the conditions necessary for the reliability of the cognitive functions involved in the acquisition of those warrants are met. One such condition involves the integrity of the (putative) concepts involved, so the needed entitlement must be reckoned to extend at least to whatever I know to be a conceptually necessary condition of that integrity. Suppose accordingly that I am an externalist and accept that an appropriate history of interaction with watery stuff is such a conceptually necessary condition for possession of *water*. Then in claiming that McKinsey (i) is warranted, I presuppose an entitlement—hence in the general sense of the term, a warrant—to take that condition to be met. A warrant for McKinsey (iv) thus becomes part of the stage-setting presupposed by my claim to possess warrant for the McKinsey premises: My route to warrant for them goes via warrant for that proposition, which is therefore ineligible to receive warrant by transmission down the extended argument. But now suppose by contrast that I regard McKinsey (ii) as articulating at best an inductively confirmed, causally necessary condition for possession of *water*. Then there has to be a fully intelligible possibility of counterexamples—cases of

thinkers who possess the concept although the appropriate history of communal interaction with the stuff is missing. So my most fundamental grounds for crediting thinkers (myself included) with grasp of the concept will be independent of such interaction—and it will be McKinsey (ii), rather than McKinsey (i), that will be in jeopardy in cases where McKinsey (iv) is false. The fourth condition of the disjunctive template (holding that conditions unsuited to the reliable operation of the cognitive faculties involved in obtaining warrant for the premises of a targeted argument would obtain if its conclusion were false—and specifically, that concepts ingredient in the content of those premises would be compromised) is therefore not satisfied; and accordingly, the template finds no problem with the transmission of warrant in the envisaged case. Rather, augmentation of the item of self-knowledge expressed by McKinsey (i) with empirical grounds for McKinsey (ii) can provide transmissible empirical support for the prediction of McKinsey (iv) in just the straightforward fashion in which, intuitively, it should.

In brief, the key difference for the bearing of the disjunctive template in the two cases turns on the question, what would follow if McKinsey (iv) proved false? If McKinsey (ii) is accepted as a conceptual necessity, there is no option but to conclude that I would lack a concept presupposed by my endorsement of McKinsey (i)—and hence that, at least in the relevant locality, a condition on the reliability of my ordinary powers of self-knowledge would be abrogated. So the fourth condition of the template kicks in. If McKinsey (ii) is accepted as an inductive generalization, on the other hand, then should McKinsey (iv) prove false, the entitlement to presuppose the good-standing of all relevant concepts would override the (merely) inductive evidence for McKinsey (ii) and it would be discarded. So, in this case the fourth condition of the template is unsatisfied and the template does not apply.¹⁰

Notes

1. The initiation and locus classicus of the modern discussion is, of course, Dretske (1970). But he does not make the distinction between transmission and closure, and the subsequent literature has largely followed him in this.

2. I borrow, of course, from Dretske 1970.

3. For example, in a paper read at the 2001 Rutgers Epistemology Conference. See Wright, forthcoming.

4. In earlier work (for instance, Wright 2000a), I had a precursor of this condition which involved *C*'s incompatibility not with the reliable operation of the

cognitive capacities involved in generating the warrant for *A* but with *A* itself. This leads, as Brian McLaughlin observed (just take *C* as not-*A*) to the template's blocking transmission—no matter what *B* is—of any noninferential warrant for *A* which one could fully convincingly *seem* to have consistently with *A*'s falsity.

5. This is of course a *closure* step.

6. See my "Cogency and Question-Begging" (2000a).

7. See Davies 1998, 2000a, and his contribution to the present volume.

8. See Boghossian 1997.

9. It should be superfluous to remark that no case of transmission failure exemplifying the disjunctive template presents a counterexample to closure. On the contrary, it is built in to the diagnosis of the transmission failure involved that they do not. The diagnosis is precisely that a *prior* warrant to accept the relevant propositions *B*, appreciable independently of their entailment by the relevant propositions *A*, is a necessary condition for possession of the relevant kinds of warrant for the latter. So there will be warrant for the conclusions of the relevant arguments whenever there is (that kind of) warrant for their premises.

10. This paper summarizes and develops my remarks at an APA Central Divisional Meetings Symposium on Externalism, Self-knowledge, and Skepticism held at Minneapolis on 6 May 2001. I am grateful for the critical comments of the other symposiasts, Martin Davies and Brian McLaughlin, and of the chairman, Michael McKinsey. The paper has been written during my tenure of a Leverhulme Research Professorship, and I gratefully acknowledge the support of the Leverhulme Trust.

This page intentionally left blank