# 10 Closure and Transmission Again

*Crispin Wright*

## 1 Closure: Dynamic or Static?

In the standard vernacular of logicians, a property of propositions is said to be closed under logical consequence if and only if it applies to every proposition entailed by any proposition to which it applies. Examples are truth, necessity, possibility and likelihood; counterexamples are falsity, impossibility and exceptionality. A further set of counterexamples are epistemic: being known, being believed, being regarded as certain. The latter counterexamples arise if for no other reason than because we are not logically far-sighted enough explicitly to acknowledge every consequence of what we know, believe or are certain about.

But one might yet be tempted to suppose that every consequence of what we know is at least know*able*—indeed that every consequence of what is knowable may likewise be known. When it is pointed out that logical consequence itself may not always be a decidable relationship, one may then retreat to the thought that every *recognisable* consequence of anything knowable is likewise knowable, that is, to

> *Knowability Closure*
>
> If P is knowable and may be known to entail Q, then Q also is knowable.

Knowability Closure is an epistemic closure principle pretty much in keeping with the logicians' standard use of "closure" — it departs only in replacing the reference to logical consequence by one to recognisable logical consequence.

However recent epistemology has departed significantly further from the logicians' standard use of "closure". Consider Matthew Jope's formulation of what he calls

> *Intuitive Closure*
>
> If S knows that P, and S competently deduces Q from P, thereby coming to believe that Q on the basis of the competent deduction while retaining their knowledge that P, then S knows that Q.[1]

DOI: 10.4324/9781003104766-11

Intuitive Closure is a clone of what Duncan Pritchard calls

> *Competent Deduction Closure*
>
> If S knows that P, and S competently deduces from P that Q, thereby forming a belief that Q on this basis while retaining her knowledge that P, then S knows that Q.[2]

These formulations add to the idea of Q's being a recognisable logical consequence of P the stipulation that its being so is actually competently deductively recognised by the subject, and then add more: that if S bases a belief in Q on that fact then—assuming knowledge of P throughout—they will then enjoy knowledge of Q. The closure of knowability across knowable entailment has thus become a principle about knowledge *acquisition*.

It is easy to miss the difference. Knowability Closure above says nothing about *how* or *why* Q may be known. It merely stipulates that the knowable consequences of known propositions should be knowable *somehow*. It says nothing about the process whereby they may be known or on what knowledge of them may be properly based. It is therefore quite consistent with the existence of cases where although P entails Q and both may be knowable, knowledge of Q cannot rationally be based on knowledge of P but must be independent of it (at least if knowledge that P is envisaged as acquired in a particular way).

This is arguably as it should be. Closure principles generally say nothing about the grounds for the entailed proposition's possessing the property in question. They are not inheritance principles. The consequences of necessary propositions are, necessarily, one and all necessary. But there is no basis in that thought for the idea that when a necessary P entails Q, Q's necessity is somehow sourced in that fact. That is a further issue. Closure may or may not be indicative of grounding.

So why does Jope count as "intuitive" a principle incorporating the additional strength of his proposal? Here is what he says:

> What makes *Intuitive Closure* so intuitive is the thought that deduction is a paradigm way of growing one's knowledge base. Williamson argues from the intuition that "deduction is a way of extending one's knowledge" to the intuitive closure principle (2000, 117). Echoing Williamson, John Hawthorne argues that "The core idea behind closure is that we can add to what we know by performing deductions on what we already know" (2005, 29). Likewise, in trying to formulate a satisfying closure principle, Steven Luper takes us to be trying to capture the intuition that "we can extend our knowledge by recognizing, and accepting thereby, things that follow from something that we know" (Luper, 2016). And in a similar spirit, Duncan Pritchard articulates the force of the intuitiveness of

> intuitive closure by asking "How could one draw a competent deduction from one's knowledge … without thereby coming to know the deduced conclusion?"[3,4]

This is a distinguished list of authorities. But the quoted thoughts of Williamson, Hawthorne, and Luper converge on the idea, merely, that deduction is (normally) *a* way of extending knowledge—which no one sane will dispute—not that each and every deduction is potentially at the service of knowledge extension. Pritchard, for his part, merely asks what is presumably intended as a rhetorical question.

Jope continues:

> One thing that is interesting to note here is that, if we take these remarks at face value and understand closure as a dynamic, knowledge-extension principle, the line between closure and transmission starts to blur.[5]

That blurring, in my judgement, is not an upshot that we should welcome. There are two reasons why we should want the distinction between Closure and Transmission to be robust.

First, while epistemologists are and should be properly concerned with the formulation and refinement of principles which constrain knowledge—or more generally—epistemic warrant *acquisition*, they also have matching reasons to be concerned with principles which properly constrain the *management* of our beliefs, with their co-ordination and revision. And static closure principles have a strong case for inclusion in the latter category. It is, on the face of it, only a static closure principle, like Knowability Closure, to which the standard closure sceptical argument need appeal: should it be indeed beyond all knowledge that one is not an envatted brain in the usual scenario, it is only static closure that would be needed to draw the sceptical conclusion that any proposition that entails that one is not so afflicted is likewise beyond all knowledge. More generally, whatever the trend in the recent literature, there are good motives to explore non-dynamic principles which concern knowledge *architecture* and knowledge *commitment*. When does knowledge of some propositions ensure the possibility of knowledge of others? When must you know certain propositions if you are to (claim to) know others?

When Dretske (1970) and following him Nozick (1981), motivated by the thought that knowledge-conferring processes should, as such, be counterfactually sensitive to variation in the fact known, proposed that Closure actually fails, the suggestion was generally received as shocking. What shocked was the excuse seemingly offered for apparently egregiously irrational professions like "I know I have hands but no-one can know whether they might not be a handless brain in a vat", which,

according to Dretske's and Nozick's theories, ought not to have seemed irrational. Readers will have to consult their own intellectual phenomenology, but I would suggest that the sense of irrationality provoked by such a claim has much less to do with the idea that in claiming knowledge that they have hands, the speaker thereby claims to be in position deductively to advance to knowledge that they are not a brain-in-a-vat than it has with the idea that in claiming knowledge their having hands, they commit themselves to the possibility of knowing that they are not a brain-in-a-vat, which accordingly cannot simultaneously rationally be disclaimed.

The second reason why we should wish to retain a non-dynamic understanding of Closure, contrasted with Transmission, is this. Suppose, as I will in this chapter once again argue, that Transmission-failure is a real phenomenon. Suppose in particular that there are cases where knowledge is achievable both that P, and that P entails Q, but where there is no possibility of rationally *basing* knowledge of Q on recognition of those facts. Then there will be a question whether—if such a case is not simply a case where Q is unknown: in other words, a straight example of closure failure as Dretske and Nozick anticipated—it is nevertheless a condition on such a case arising that Q be known independently. If that is so, then that will be an argument for sustaining a static closure principle while rejecting any unrestricted dynamic principles such as Intuitive Closure and Competent Deduction Closure.

What is unfortunate about an exclusive focus on dynamic closure principles is that it encourages oversight of that dialectic. Generally: when Closure (for whatever epistemic operator) is formulated dynamically, as in Intuitive Closure and Competent Deduction Closure, there is of course, as Jope observes, foreseeable difficulty in explaining how it can hold when Transmission fails. But that is because dynamic closure is already, in effect, a principle of unrestricted transmission. If Transmission does indeed fail in certain cases, the question is whether—and if so why—some form of static closure principle still holds, or whether none does. The tendency exemplified by Jope and the authors he cites is liable to make that important question inaudible.

In what follows, I will review the case for rejecting unrestricted Transmission before returning to the question of the status of Static Closure principles. First we need to be clearer about this question:

## 2 What Is it for an Argument to Be Transmissive?

What is it for a valid argument to be transmissive of epistemic warrant for its premises to its conclusion? I have explored various accounts in previous work[6] but here I think it will be helpful to focus on one especially clear kind of diagnosis.

It will call for some kit. First we need what I think, since the work of Pollock and others, is a standard notion of an *undermining*[7] defeater for a belief. Suppose we are in the business of accumulating evidence on the question whether or not a proposition P obtains. Let E be the evidence so far gathered and suppose we incline to regard E as defeasibly supportive of P. Let U state a potential underminer of E's support for P: that is, (we are of the opinion that) if U is true, then E should no longer be regarded as supportive of P.

For example, let P be that you have contracted COVID 19, and let E be the positive outcome of a lateral flow test that you have just taken using a kit purportedly supplied by a well-known pharmaceuticals manufacturer based in San Diego, CA. We can count it as an additional part of your evidence that the kit is advertised at approximately 98% positive accuracy—false positives occur on average only in 1 of 50 trials. Now suppose you learn—this is U—that the Wolverhampton depot in which the kits are stored for distribution in the UK has inadvertently taken in a large number of fake kits. Presumably, this information should rationally be received as to some degree compromising E's support for P.

Note that whether something is an underminer of the significance of a given item of evidence may vary with collateral information. It may only be in the context of background theoretical knowledge that a given kind of evidence for a certain belief loses its supportive force. Up on the Cuillin Ridge in the Isle of Skye in thick mist, I get a steady compass signal (E) indicative, in conjunction with the map, that I am indeed on track and should continue on in the same direction. Someone says, "You know the rock around here is predominantly gabbro?" While I ought to regard this as undermining the seemingly encouraging compass reading, I am ignorant of the magnetic properties of gabbro. So I merely reply "Yes", and carry on trusting my compass undaunted. We can prescind from such information-relativity in the notion of an underminer in what follows.

Second, we need a notion of *open-mindedness*. Your open-mindedness about P, in the sense that concerns us here, requires

i your having no opinion whether or not P;
ii your not being *agnostic* about P in the sense of "agnostic" frequently associated with the stance of religious agnosticism, that is a stance which holds e.g. that the existence of God is a matter on which no rational subject can in any circumstances can consider that they have a knowledgeable or even a justified view. Your open-mindedness about P involves readiness to accept or reject P, should appropriate evidence turn up but also, as we shall here understand it,
iii your being of the opinion that *so far* you have no basis for a view, but not ruling out the possibility of new information that might mandate taking one.

With these two notions—open-mindedness and undermining defeat of evidence—in place, we can elicit an evident but interesting corollary. Suppose you are presented with purported evidence E, for P, and suppose you are open-minded about whether something you recognise as a specific potential underminer, U, for E obtains. Note that we are not supposing that you are *sceptical* about U—doubtful that U is true—but merely that you consider that you so far have no basis for a view about whether or not U is true. Well, if U does obtain, then, as you recognise, E should not be regarded as supportive of P. So if you are open-minded about U, you ought rationally to be open minded about whether or not E should be regarded as supportive of P.

More explicitly, here is the reasoning that underlies that thought. U, if true, is an underminer of E's evidential force for P. So if E is supportive of P, U is false. So if your present state is one of open mindedness whether U is false or not, it cannot be that you are in epistemic position to regard E as supportive of P.

There is an implicit appeal here to closure for whatever relevant epistemic notion is supposed to be at odds with open-mindedness. But that seems perfectly intuitive. Consider an example. You are running a schoolroom physics experiment in which it is important that there be no leakage of pressure from the apparatus. If pressure is lost, = U, the significance of the findings for the hypothesis under test, = P, will be severely compromised. Suppose you are open-minded about whether pressure *was* lost; members of the class attached the clamps, you haven't checked their work and reckon it quite possible it was slapdash. But you run the experiment, take the final readings, = E, and announce that they are indeed confirmatory of P. Can you rationally profess to have confirmed P although simultaneously open minded-about whether pressure was indeed lost?

We have introduced the relevant options here in relation to cases where what is at issue is the suasive force of defeasible evidence. But the same basic framework applies equally to cases where P is based on some form of (putatively) direct cognition: perception, episodic memory, or (perhaps) non-inferential intellectual intuition, for example, where reliance on defeasible evidence is not involved. Even so, cognitive faculties of this kind have conditions of effective operation. And a claim based on the presumed operation of such a faculty may be undermined by finding that these conditions were compromised in one way or another. Perhaps the conditions under which the faculty was operating are unsuitable for its effective function. (Suppose, for instance, you are in Barn Façade County.) Or perhaps the faculty itself was suffering some kind of internal malfunction. (Perhaps you have taken an hallucinogenic substance of some kind.) Then, once again, you cannot rationally advance a claim, P, based on the presumed effective operation of such a faculty, and at the same time profess open-mindedness about whether some condition,

U, obtains which you acknowledge to be a potential underminer of the faculty—a condition such that, if it obtains, the effective functioning of the faculty concerned will indeed be likely to be compromised.

We still need two more notions. Where U is any underminer of the support of evidence E for P, let us call the negation of U an *Authenticity Condition* for E with respect to P. Likewise where U is an underminer of the capacity of faculty F to reliably determine whether P, let the negation of U be an Authenticity Condition for F with respect to P.

Finally let us say that an argument {P ⇒ Q} is potentially *cogent* for a rational thinker who is presented with evidence E for P, or who achieves prospective verification of P by employing faculty F, just if acknowledgement of the evidential force of E for P, or the competence of F in the circumstances to determine whether P,[8] is rationally consistent with open-mindedness about Q. Intuitively, a cogent argument is one where acceptance of the grounds offered for its premises is consistent with anterior open-mindedness about its conclusion. Within the class of deductively valid arguments, it is only the cogent arguments, thus characterised, that are at the service of extending our knowledge.

So we come finally to the crux: is there any other kind? Are there valid arguments for whose premises there are probative grounds, but which are nevertheless not cogent for a given thinker who is antecedently open-minded about their conclusion?

Well yes, indeed there are! We now have all we need to articulate a template for one kind of Transmission failure.[9] We merely have to construct a valid argument, {P ⇒ Q} where Q is an authenticity condition for the particular evidence, E, offered for P, or for the competence of F to determine whether P. In order for the argument to be potentially cogent, it has to be rationally possible to be open-minded about Q consistently with appreciation of the probative force of E for P. If Q is an authenticity condition for E with respect to P, that is exactly what is not rationally possible.

The non-transmissiveness of an argument is accordingly not a feature purely of the contents of the premises and conclusion concerned but is relative to the informational context: it is relative to the character of the particular epistemic pedigree of the premises. Non-transmissive but valid arguments for a given thinker are arguments where, such is the thinker's information, a commitment to the conclusion, whether or not based on independent cognitive achievement, is a rational precondition of their accrediting the evidence for the premises (or the cognitive process involved in recognising them as true) as persuasive of their truth. You cannot consider yourself to have learned anything by following such an argument if prior persuasion of the conclusion is a rational precondition of accepting the premises on the particular grounds proffered.

Here are a couple of quotidian examples:

### *2.1 Chiff-Chaffs and Willow-Warblers*

Suppose—to adapt an example of Austin's—that I am a less than expert but not totally incompetent bird watcher whose method of deciding whether or not a bird is a willow warbler is simply to judge on the basis of size, colour, movement and overall appearance—*provided* I am already confident it is not a chiff-chaff. If, on the other hand, I regard that issue as open, then, the two species being so similar, I have no method except to ask an expert. Let the region be such that there are vanishingly few chiff-chaffs in the local bird population—indeed, no warbler-family birds of any kind other than willow-warblers. And suppose I believe this. Then the following argument: —

E: That bird's size, colour, movement and general *gestalt* suggest it is a willow-warbler
P: That bird is a willow warbler
Q: That bird is not a chiff-chaff

—is non-transmissive for me, though not for an expert birder who can distinguish the two species by sight. That is because, given my background information and limitations, I can rationally treat E as probative of P only if I am antecedently confident of Q. Note that the situation is not that I get a warrant for P in any case which then, as it were, gets stuck on the wrong side of the entailment from P to Q. Rather it is only in a context where I am antecedently confident of Q that I can rationally consider that I have good grounds for P. I cannot rationally appeal to the argument to enhance my confidence in Q.

Another example

### *2.2 Twins*

Jessica and Jocelyn are identical twins with a disconcerting habit of matching their daily attire, hairstyle, etc. Both are well known to me but not so well that, seeing either by herself, I can be confident which of them it is. Consider this argument:

E: That girl looks exactly like Jessica
P: That girl is Jessica
Q: That girl is not Jocelyn

Let the context be one where I am antecedently confident of Q: I know that Jocelyn is out of town to attend a campus fly-back. Then, seeing a girl approaching who looks exactly like Jessica—so I have E—I am fully warranted in concluding P. But I cannot then rationally refer this warrant across the entailment to Q to conclude that she is not Jocelyn. Rather

again, it is only in the context of my antecedent knowledge about Jocelyn's whereabouts that I can consider E supportive of P in the first place.

One more case of a slightly different structure:

### *2.3 Town Clock*

I am at the Urbino summer school in Epistemology sitting out in the town square during the mid-morning break and wondering if I have time for another excellent Italian coffee before the 11.00 am session. I look across at the town hall clock and note that it says 10.51. So there is just about time. Can I trust it to be accurate enough? (I know it is functioning—it chimed a few minutes ago.) I reason as follows:

E: The clock says 10.51
P: So the time is 10.51
Q: So (from E and P together) The clock is accurate on this occasion.

Thus (idiotically) reassured, I catch the waiter's eye … …

These examples ought to be persuasive that rational belief management sometimes requires recognising that the order of the epistemic dependencies among one's beliefs and suppositions inverts the order of their logical dependencies. But notice that we have taken the issues into the territory of rational belief management. Our question has been, in effect, can I always rationally transmit whatever grounds I have for certain beliefs to propositions that I recognise to be entailed by them? The answer manifestly is: No—not when the latter are authenticity conditions for the probative force of the grounds concerned for the former.

But does this point abrade with dynamic closure principles for knowledge?

## 3 Does transmission failure, as now schematised in section 2, count against the validity of intuitive closure?

Jope's *Intuitive Closure*, recall, runs as follows:

> If S knows that P, and S competently deduces Q from P, thereby coming to believe that Q on the basis of the competent deduction while retaining their knowledge that P, then S knows that Q.

Closure "Dynamists"—Jope, Pritchard, Hawthorne, Williamson et al.—might argue that there is no tension between this kind of principle and the gist of the previous section. After all, the examples offered illustrate how a thinker's open-mindedness about the conclusions of certain valid arguments, $\{P \Rightarrow Q\}$,—a precondition of their potential cogency—ought

rationally to have the effect of disabling a prima facie warrant for their premises. And in that case, the thinker won't know that P in the first place, so that the antecedent of *Intuitive Closure* will not be satisfied.

But not so fast: matters are more complex. First, we must be mindful that *basing*, as a real psychological phenomenon, can be rationally inappropriate: one can base a belief on considerations that are poorly selected to justify it. Biases, prejudices and favouritism offer a legion of examples where a thinker's belief is based on considerations that are rationally inappropriate. Second, we ought to want to impose a rationality condition on knowledge in general. A belief should not count as knowledgeable, even if gifted by the circumstances of its formation—the method and prevailing conditions—in such a way as to ensure its truth, if the thinker is nevertheless irrational to hold it. (If some of your prejudices turn out to be nomically true, that doesn't place the beliefs they have driven you to hold in epistemic good standing.)

The question is accordingly whether there are circumstances of the structure depicted by *Intuitive Closure* in which basing the belief in Q on P and knowledge of the deduction would result in a true—since the consequences of known propositions are true—but irrationally based belief. And the answer will depend on how relaxed we are about the degree of epistemic scrutiny required to know that P. Consider *Town Clock* again. Could I be said to know what time it is by looking at the Urbino clock? Ordinarily we should not scruple to allow that I could. No doubt that's because it's common knowledge that clocks are designed to tell the time, that municipal clocks are generally maintained so as to help the citizens rather than provoke confusion, etc. Generally it would seem in most contexts overly fastidious to insist on a background check on the clock in question before regarding its readings as knowledge-conducive. So allow that I can know what time it just by looking at the clock. Manifestly I can also know what the clock says just by looking at the clock. So, by *Intuitive Closure*, I can know the conjunction—that the clock says that T and that T, and hence that the clock is accurate on this occasion—just by looking at the clock and basing my belief in the conjunction just on a conjunction-introduction step. But that is manifestly an irrational basis for confidence in the on-the-occasion accuracy of the clock.

In general, *Intuitive Closure* will be in trouble whenever Q is an authenticity condition for the relevant method for acquiring knowledge of P of such a kind that ordinary standards, or a preferred philosophical account, of what it takes to get knowledge of P requires no independent scrutiny of Q which is, rather, permissibly "taken for granted", as we say, if considered at all—a "prop" in the sense of fn. 9 above. In neither case would basing the belief in Q on the deduction amount to basing it on anything that corroborates it. So: where knowledge is required to be rational, *Intuitive Closure* as formulated is open to counterexamples and accordingly is false.[10]

What *might* be true would be something along these lines:

> *Intuitive Closure**
>
> If S knows that P *by some method for which Q is not an authenticity condition*, and S competently deduces Q from P, thereby coming to believe that Q on the basis of the competent deduction while retaining their knowledge that P, then S knows that Q.

But I shall not explore further here whether further tweaks are needed.

## 4 Transmission Failure—Some Controversial Cases

Let us now briefly review, in the light of the Transmission-failure template outlined in Section 2, four somewhat notorious and much discussed arguments: G.E. Moore's[11] (widely ridiculed until quite recently) 'Proof of an external world', Michael McKinsey's[12] apparent proof of the inconsistency of semantic externalism with privileged psychological self-knowledge, Hilary Putnam's[13] (generally disbelieved) proof that we are not brains-in-a-vat, and Fred Dretske's[14] fraudulent zebras.

First up is

### *4.1 Moore*

The status of Moore's 'Proof' appears to turn on what we write in for E. *This* reasoning, at least, is arguably non-transmissive:

E: My experience is in all respects as if there is a hand in front of my face
P: Here is a hand.

So, (since hands are material objects existing in space) —

Q: There is an external material world.

Q is manifestly an authenticity condition for the probative relevance of E for P. If, bizarrely, I am open-minded about the existence of the external world, I ought to be open-minded about the evidential relevance of my experience for claims about it.

There are two ways this account of the shortcoming in this version of *Moore* may be challenged. One is the perceptual Dogmatism stoutly defended by Jim Pryor[15] On Pryor's account, Q while indeed an anti-underminer for E *vis-à-vis* P, is *not* such that antecedent open-mindedness about it must rationally require open-mindedness about the evidential force of E for P. What is true is that an antecedent *doubt* about Q—a tendency to disbelief of Q—defeats the evidential force of E for P. But mere open-mindedness should not. Rather sense experience

is properly regarded as default (though defeasibly)—in post- *Philosophical Investigations* jargon, *criterially*—evidentially supportive for claims about the external world and may thus, absent defeating evidence or spoiling background considerations, be rationally responded to by the formation of appropriate external worldly beliefs. I shall not further discuss perceptual Dogmatism here.[16]

A second response, currently somewhat in the ascendant, is to dispute the very epistemic architecture of the proof when it is configured as in the 3-step form above. On this kind of response, it is wrong to interpret Moore's train of thought as beginning with evidence conceived as internal to consciousness. Rather the sensory states on which perceptual justification is based should be conceived of as *wide*—as world-involving. When Moore holds up his hand in front of his face, the epistemic ground for his claim that "Here is a hand" is directly constituted by his sensory awareness of the hand, an intake, in John McDowell's phrase, of an aspect of "the layout of reality".[17] In effect, therefore, the proper way of representing Moore's proof is rather this:

### *4.2 Moore Lite*

E/P: Here is a hand. So
Q: There is an external material world

When the architecture of the proof is so conceived there is, of course, no transmission failure. But it is a further question whether, as seems to be widely supposed, this adjustment—the resort to a factive notion of perceptual warrant—can be marshalled against perceptual scepticism without, in effect falling back on reasoning for which Transmission-failure is once again an issue. The point can best be bought out in a "disjunctivist" setting. Consider this argument:

### *4.3 Ante-Moore*

E: The phenomenal character of my subjective experience is consistent both with its amounting to perceptual awareness of a hand in front of my face and with its being the product of a sustained hallucination.
P: I am perceptually aware of a hand in front of my face
Q: Here is a hand

Here an endorsement of E is consistent with acknowledging, with metaphysical disjunctivism, that there is no one kind of experiential state of which both veridical perception and hallucination are species. Rather the one is, metaphysically essentially, a form of relational engagement with matters external while the other is but a play of light and colour in phenomenal consciousness. Let that be so. Still, the point hasn't gone

away that the one kind of experience can in principle be subjectively indistinguishable from the other. So if we represent the ultimate justification of the premise in *Moore Lite*—"Here is a hand"—as correctly captured in its role as conclusion in *Ante-Moore*, then the issues about transmission failure are now going arise in the process of getting as far the premise of Moore's Proof, not in the drawing of its metaphysically heavyweight conclusion.

In general, it seems that any kind of reasoning that starts with a basis of evidence given by phenomenal consciousness and purports to underwrite conclusions about the very existence of a material reality is going—*pace* Pryor's dogmatism—to be open to a challenge of Transmission-failure unless some case is made that that is a misconceived starting-place: that perceptual justification *begins* with factive states of worldly awareness. The sceptic will then gleefully rejoin that their challenge is in effect to justify the claim that we can legitimately start with the assumption that we do indeed enjoy such factive states. At this point the issue with Moore's proof is not one of whether the reasoning is warrant-transmissive but whether we are entitled—and if so what entitles us—to lay claim to its premise.[18]

Now to *McKinsey*:

Michael Mckinsey's much discussed 'paradox' purports to elicit a contradiction between the status of ordinary, 'privileged' self-knowledge of one's intentional attitudes and one form of semantic externalism. It is simply formulated:

P(1): I believe that water is wet
P(2): An agent has the concept of water only if they, or others of their speech community have, historically, interacted with water. Hence
Q: Members of my speech community have, historically, interacted with water.

The paradox is then that ordinarily effortless and immediate self-knowledge of one's basic intentional states can deliver knowledge of P(1). So if a priori philosophical reflection can indeed corroborate a content externalism that entails P(2), viz. that it is only if the concept of water picks out a real physical kind with which we have had historical interactions that it is a real truth-conditionally contributive concept, then it would appear to follow that by combining quotidian armchair psychological self-knowledge with some subtle armchair philosophy, I can deliver knowledge of a contingency of history—pretty good going from the armchair!

If you share my hunch, you will suspect that Mckinsey's argumentation is somehow non-transmissive—that, on its assumptions, it somehow merely elicits a presupposition of the warrant for its premises, rather than furnishing a genuinely novel reason for believing its conclusion. But

is that so? There are two obstacles in the way of getting a clear view of the matter.

First, a mechanical application of the template for Transmission-failure proposed in Section 2 would bid us ask whether open-mindedness about *McKinsey*'s conclusion would rationally enforce open mindedness about the sufficiency of the warrants for its premises. But what *are* the warrants for its premises? Well, in the case of P(2), the warrant is given by the philosophical considerations that are supposed to mandate the relevant content-externalism. One might, of course, have one's doubts about how persuasive those considerations really are. But it doesn't seem that open-mindedness about whether members of my speech community have ever historically interacted with water should enjoin open mindedness about the cogency of such philosophical argumentation *as a species*,—not in the way that open-mindedness about the very existence of an external world ought to entrain open mindedness about the probative force of phenomenal experience for claims about material objects. As for P(1), well what in any case *is* our warrant for sincere self-ascriptions of basic intentional attitudes? I think we have no clarity about the cognitive machinery underlying avowals of the kind expressed by P(1), and consequently no clarity about the conditions for its effective operation, or what circumstances might undermine it.

The second awkwardness about trying to bring *McKinsey* under the Transmission-failure template is that open-mindedness about the conclusion—"Members of my speech community have, historically, interacted with water"—is *itself* an attitude that embeds the concept of water and thereby, in the presence of P(2), entails the integrity of that concept which, from an externalist point of view, is itself already a commitment to the conclusion. You cannot rationally be open-minded about Q in this case without denying P(2)—so trivially you should wonder about the adequacy of the grounds for P(2). But the kind of commitment to so wondering which the test is looking for is not supposed to go through a doubt about the truth-value of a premise, as opposed to a doubt about the adequacy of the type of grounds proffered for it.

We could just give up on the suspicion that *McKinsey* is non-transmissive. But I believe that would be premature.

Here is another approach. Content-externalist views come with the possibility of certain kinds of *illusion* of content. Most of us would, for example, accept an externalist view of singular demonstrative content. When we use demonstratives, real external worldly semantic relations are required before we can succeed in thinking anything with a determinate truth condition. Suppose I am a Victorian lepidopterist on an Amazonian expedition, looking for rare and exotic species. Unmindful of the rather delicious stew of the local fungi I lunched on, I seemingly catch sight of a large specimen, brilliantly patterned in orange, purple and aquamarine, and think to myself, "That is a hitherto undocumented

species and will make me famous". I do not think anything true. But nor is there any determinately false singular thought which is exactly what I think—the nearest determinate false thought will be an existential thought: that there is a brilliantly coloured butterfly over there, etc. In so far as the content I intentionally think purports to be a *demonstrative* content, it is an illusory content.

In an influential early discussion of these matters,[19] Paul Boghossian introduced the example of Dry Earth. Dry Earth seems to its inhabitants (intrinsic duplicates of us) just as Earth seems to us: it seems to them that there are rivers, lakes and rain, and a clear, odourless, potable liquid flowing from their domestic taps. But this is all illusory. There's no such substance anywhere in their environment. What goes wrong when a Dry Earthling runs Mckinsey's argument? Well clearly, the premise, P(1), is not true: the embedded content, *that water is wet*, is, by externalist lights, an illusory content for the Dry Earthling, so the tokening of 'water', no less than the tokening of 'that' in the lepidopterist's thought, is divested of determinate truth-conditionally contributive content. The *Mckinsey* Q is an authenticity condition not for the good standing of the process whereby P(1) is seemingly confirmed but for there being anything determinate to be confirmed in the first place.

That seems nearer the mark. But there is still the awkwardness that *McKinsey* Q, as formulated, presupposes the good standing of the embedded content in P, so that the relevant issue is silenced just by open-mindedness about Q. I think there is no way around this except to countenance the notion of what is, by externalist lights, a *defective* concept that can nevertheless feature in *ersatz* attitudinal states that, while contentually defective, are nevertheless genuinely explanatory of thought and action in routine ways. (It is, after all, not as if all is dark within the Victorian lepidopterist's head: suppose, e.g., he gets out his net and takes — just for the purposes of section 4.4 — a swipe at what is in fact thin air. How to explain his action?) So, let us adopt the convention of i*talicising* a concept expression to indicate that it is to convey a 'concept' which, in the good case, where external presuppositions of its purported content are met, will indeed make whatever kind of semantic contribution externalism proposes, but in the bad case will contribute only to an illusion of determinate thought. And let us say that a 'concept' of this kind is *validated* just in case the relevant external conditions for its determinacy of content are indeed met. Then we can re-run a Mckinseyish argument like this:

### *4.4 Ante-McKinsey*

P($1^{minus}$): I believe that *water* is wet
P(1): I believe that water is wet
P(2): An agent has the concept of water only if they, or others of their speech community have, historically, interacted with water. Hence

Q: Members of my speech community have, historically, interacted with water.

Now we have in effect equipped ourselves with the resources to identify an authenticity condition for McKinsey's P(1). It is:

Q*: Members of my speech community have, historically, interacted with a substance that validates *water*.

Note that, given that the only substance that can validate water is … water (!), Q* is actually equivalent to Q.

Open-mindedness about Q* rationally enjoins open-mindedness about the transition from $P(1^{minus})$ to P(1). If I am open-minded about the possibility that *water*, in my thought, is by externalist lights, contentually illusory, then I should be open-minded about whether I can justifiably advance from the piece of self-knowledge expressed by $P(1^{minus})$ to McKinsey's original P(1).

If this analysis of the issues is accepted, the suspicion that *McKinsey* merely concludes in a proposition that the justifiability of its premises presupposes, rather than generating a novel reason for accepting that proposition, is confirmed. But the explanation is not exactly that the actual reasoning McKinsey gave is non-transmissive. *McKinsey* itself, Like *Moore-Lite*, is transmissive. Nevertheless the impact of the relevant kind of content-externalism is to complicate the justificational architecture of its premise P(1) in such a way that (i) the real item of pure self-knowledge, strictly so regarded, involved is rather $P(1^{minus})$, and that (ii) there is an authenticity condition for the transition from that to *McKinsey*'s P(1) — to wit: Q*—which is equivalent to the original conclusion Q.

### *4.5 Putnam*

Putnam's 'Proof', like Mckinsey's, is fuelled by somewhat inchoate content-externalist assumptions. He envisages a scenario in which creatures whose subjective mental lives are qualitatively indistinguishable from ours, and whose language is lexically and grammatically indistinguishable from English, have in fact always been brains in a vat, tended by computers and automatic machinery; moreover this description of their situation is actually a complete inventory of their world. Putnam takes it that, in this predicament, these thinkers are so causally isolated from their real environment that, on externalist assumptions, they could not so much as have the concept of *brain-in-a-vat* and thereby could not mean *brain-in-a-vat* by their Anglo-form expression, "brain-in-a-vat", or by any other. However we, in regular English, mean exactly *brain-in-a-vat* by "brain-in-a-vat". It follows that we are not in the described predicament—that we are not brains in-a-vat.

Putnam's argument makes assumptions that compromise its possible epistemological significance. His scenario has features—for instance, the brains' historically complete causal insulation from brains or vats—which are required to get the background content-externalism to engage with it but which someone who wanted to fashion a brain-in-a-vat scenario for sceptical purposes could easily dance around, e.g. by working with the alternative scenario of overnight envatment by a mad scientist instead. In fairness, however, Putnam himself seems to have primarily conceived the argument not as a riposte to scepticism but as a challenge to a core idea of metaphysical realism—roughly, the notion that if our epistemic entanglement with the external world generates any true beliefs, that it does so is matter of pure contingency. I won't comment further here on its effectiveness in that direction.[20] Our question is the limited one whether when one reasons as below, and assuming the premises are warranted, one generates a warrant for accepting the conclusion:

*Putnam*:

P(1): In the language I speak, 'brain-in-a-vat' means *brain-in-a-vat* (My language is disquotational)

P(2): In the language of brains-in-a vat in the described predicament, "brain-in-a-vat" could not mean *brain-in-a vat* (by the absence of the causal connections which according to content-externalism would be required by such a content). So

Q: I am not such a brain-in-a-vat

This is indeed suspiciously reminiscent of *McKinsey*. But let's see. Here we have two independently supported premises conjunctively entailing the stated conclusion. The question about transmissiveness is therefore whether a rational thinker who is open-minded about—so, presumably, haunted by the possibility of—the conclusion's being true should accordingly be open-minded about the acceptability of either premise on the grounds offered for them. It can seem obvious—and has, to a number of commentators—that in the presence of P(2), P(1) is question-begging; or, as we can more felicitously now express the charge, that anterior open-mindedness about Q should, if one accepts P(2), rationally require open-mindedness about P(1). After all, if you were a brain-in-a-vat, you would by P(2) have no means of expressing the concept *brain-in-vat*.

That, though, second thoughts should rapidly teach, is an addled train of thought. Suppose you are genuinely open-minded whether or not you are a brain-in-a-vat. Nevertheless, whether you are or not, it is presumably solid that your language, the language we are now working in—whether it is English or its vat homophone—allows for the expression of the—that is, *your*— concept *brain-in-a-vat*. After all, this is the very concept about your falling under which you are open-minded. So P(1) is unexceptionable whatever your predicament. Hence if you are rightly

persuaded by content-externalist considerations that creatures in a certain scenario would lack any expression for a concept which you can express, that is just to say that you can do something which they cannot. So you are not one of them. QED

What is true is that a brain-in-a-vat could presumably carry out token reasoning of exactly this form, and accordingly validly draw some conclusion or other. "So how do I know that that is not what I have done?" You should know better than to ask! Your question is answered by the very reflection that you are not a brain-in-a-vat.

I suggest accordingly that there is no clear-headed obstacle to the idea that Putnam's argument, unlike Mckinsey's, *is* warrant-transmissive. What exactly is its philosophical significance is, of course, quite another matter.[21]

Finally

*Zebras*:

Here is Dretske's famous example, which he originally conceived, of course, as a failure of Closure:

E: The look of the animals, in a zoo cage marked "Plains Zebra"
P: Those animals are Zebras
Q: Those animals are not mules, their coats cleverly dyed and manes coiffured to look exactly like zebras.

This impresses as an absolute paradigm of Transmission-failure. I believe it has been a primary force in persuading many theorists of the reality of the phenomenon. Surely, if you are open minded about the possibility that the zoo has gone in for artful disguise and fraudulent labelling of its animals, you have to be open-minded about the significance of the appearance presented by its exhibits, in which case you cannot yet trust that E is indeed probative for P.

This was, for a long time, my own reaction to the example. But philosophical appearances, no less than those of the animals in fraudulent zoos, can of course be deceptive. What is the logical form of the above Q? Presumably it is a negated conjunction: it is not the case that [*Those animals are mules and those animals have had their coats cleverly dyed… etc.*], so it is something of the form ~(Fx & Gx). Hence it is tantamount, in classical logic at least[22] to a disjunction of negated disjuncts:

$$\sim Fx \vee \sim Gx.$$

But now the Transmission-failure diagnosis of *Zebras* hits a snag: surely any reasoning that warrants a certain conclusion *must* also warrant the disjunction of that conclusion with *any* proposition. The following reasoning is surely transmissive, notwithstanding the inconclusiveness of E:

E: The look of the animals, in a zoo cage marked "Plains Zebra"
P: Those animals are Zebras
Q: These animals are not mules.

So how can it make a difference if we tack an arbitrary disjunct onto Q?

The very tempting answer is: it cannot. What is happening, it may be suggested, when the argument is taken in a way that incites the intuition that Dretske intended, is that, by some quirk of the pragmatics of conversation, we receive a negated conjunction as though it implied the negation of its second conjunct. "He's not a very funny comedian" would usually, for example, be heard as saying that the gentleman in question is not very funny. Which is fine if the conversational common ground includes that he is a comedian. But it seems that in Dretske's example, the same conversational habit kicks in even though it is not presumed that the animals in question are mules. We hear the conclusion of the argument as implying that the animals in question have not been cleverly disguised. And while that is indeed an authenticity condition for the transition from E to P, it is not actually entailed by the relevant Q.

So, one may now be inclined to conclude, the impression that Dretske's argument has given so many as being, if not an example of failure of closure, at least non-transmissive, is in a way correct: the argument does indeed fail to transmit justification from its premise to the conclusion that the animals in question have not been cleverly disguised. But it fails to do so because, so interpreted, it is invalid!

However, having been persuaded of this reappraisal of Dretske's example for the last decade or so since Aidan McGlynn first argued for it in a seminar at the Northern Institute of Philosophy in Aberdeen,[23] I have now[24] come to think that it too is mistaken and that first thoughts were correct: for the impression that Zebras is indeed an instance of Transmission-failure is validated by the template for the phenomenon outlined in Section 2 above.

Here's why. There is nothing relevantly objectionable about representing the conclusion of *Zebras* as of the form, ~(Fx & Gx), nor with converting that to the disjunction, ~Fx V ~Gx. Rather, very surprisingly, the thought that warrant for a proposition must constitute, ergo transmit, to warrant for any disjunction in which it is a disjunct is simply wrong. For suppose you are open minded about ~Fx V ~Gx. Then you should consider that you are in no position to exclude the possibility that that disjunction is false. But that is the possibility that each disjunct is false, so that Fx and Gx are both true. And that, in the case at hand, is the possibility that you are confronting cleverly disguised mules! But if you are open-minded about that, you had better be open minded about the probative force of your visual evidence for the claim that they are zebras.

*Zebras* can thus continue to enjoy its reputation as a masthead example of Transmission-failure.

## 5 Closure and Metaphysical "Heavyweights"

For any given cognitive project, where we aim to determine a view about P by scouting indirect evidence or bringing some direct cognitive capacity

to bear, the authenticity conditions come in three broad categories. There will be some, first, which a properly responsible execution of the project will demand that we check independently and where failure to do so will be considered negligent and to compromise the significance of the results. Which conditions come into this category will often be in part a function of opportunity, stakes and interests: an untested assumption can be negligent in one context but permissibly taken for granted in another.

Second, there will be, on the opposite side of the coin, what above I termed *props*—authenticity conditions which in context, it is considered rationally permissible to take for granted, and a check on which will be considered overly fastidious or even paranoid: demanding, for instance, wherever practicable, a character reference for any casual informant on some mundane matter, constantly checking ones watch against the Speaking Clock[25]; poring over long-term analyses of the accuracy of Met office weather forecasts before being willing to plan a picnic on the basis of tomorrow's forecast that it will be sunny and dry all day.

Third, however, will be the metaphysical "heavyweights" (Hawthorne), the structural 'Hinges', that feature in what I called, in earlier work, the I–II–III sceptical arguments illustrated in this table:

| *Domain of Enquiry* | *External World* | *Other Minds* | *The Past* | *Induction* |
|---|---|---|---|---|
| I E(vidence) | Visual and proprioceptive evidence as of a hand in front of my face | Your gashed shin and twisted ankle after falling off your bicycle | Recent excavation of a huge fossilised reptilian skeleton | All Fs in a large random cross-sample of Fs prove to be G |
| II P | Here is a hand | You are shocked and hurting | A dinosaur died here many millennia ago | All F's are G |
| III Q (Authenticity Condition) | There is an external world | There are other minds | The negation of Russell's Five Minute Hypothesis[26] | The World is inductively amenable |

In each of the four columns there is at least a strong prima facie case that antecedent open-mindedness about Q would debar a rational subject with crediting the evidence of type E, variously illustrated, with probative force for any proposition of the type illustrated by the relevant P. If that is accepted, it follows that there is no getting to know the Q-propositions by the kind of inferential route indicated by the downwards vertical path in the respective columns, and hence—at least so the Sceptical argument runs—no getting to know them at all. How might we resist this conclusion?

It can be resisted if it is possible, after all, to validate something like Dynamic Closure in full generality. If Dynamic Closure could be shown to hold, then we can rest content with ordinary standards of knowledge or justification for the P-type propositions as illustrated in the table and then appeal to Dynamic Closure to quash any inhibition about whether we can thereby know or justify the Q-propositions. Let me accordingly close by offering some reactions to Duncan's Pritchard's recent attempts[27] to save Dynamic Closure in the face of the pressure exerted by the apparent counterexamples spawned by the Transmission Failure template schematised in Section 2.

Pritchard is clear that his discussion concerns the *rational* formation of knowledge[28]—a commendable departure from the cruder forms of knowledge externalism that have been prominent in the discussion of recent decades—and that the question is accordingly how, in the awkward cases, a belief in Q can be *rationalised* on the basis of the evidence for P and the deduction—exactly what the intuition of Transmission-failure denies. His strategy is to offer different treatments for cases where the Q-proposition is not a metaphysically heavyweight hinge, taking *Zebras* as epitomising the latter type of case. Lest the reader feels that the status of that example remains controversial, notwithstanding the discussion in the previous section, I will here focus on the earlier example of *Twins*.

E: My perception of a girl approaching who looks and dresses exactly like Jessica
P: That girl is Jessica
Q: That girl is not Jocelyn.

Pritchard has two proposals for saving Dynamic closure in such a case. One is to construe the evidence concerned as *factive*, so that E becomes:

E+: My perception of Jessica approaching.

The other strategy is to amplify the relevant evidence for P so that it includes sufficient background data to warrant Q: to let my evidence for P include, for example, the collateral information that Jocelyn is out of town today.

Construing perceptual evidence as factive is, as noted earlier, the key thought of disjunctivist conceptions of perceptual justification. It is, as Pritchard acknowledges, highly controversial. I have no space here to evaluate further its promise as a weapon against perceptual scepticism.[29] The more immediate point is that it cannot plausibly generalise to *prima facie* Transmission-failure cases where the E-place is occupied by defeasible evidence that P. For example, suppose I have misplaced my car keys and am going to be late unless I find them quickly:

***5.1 Keys***

E: Catherine tells me she thinks she saw them on the shelf in the pantry
P: My keys are on the shelf in the pantry
Q: Catherine is right on this occasion.

Q is an authenticity condition for the transition from E to P: Open mindedness about Q will rationally require open mindedness about the evidential significance of E for P. But here there is no possibility of taking E to be *factive*—E can be true while P is false.

What about the alternative of reconceiving my evidence for P to include background materials corroborative of Q? Perhaps Catherine has been able to bale out my forgetfulness many times in the past. Now I do indeed have evidence for Q, so a static version of Closure is comfortable with the example. But it is hard to see how amplifying the background evidence is this way addresses the concern about *Dynamic* Closure. For now the basis for my acceptance of Q is nothing to do with its deduction from E and P but is simply Catherine's track record as a helpful St. Anthony, whereas Dynamic Closure requires a basis for accepting Q in the competent deduction. As in all relevant cases, the rationality of accepting P on the basis of E depends on prior acceptance of (the likelihood) of Q. Dynamic Closure simply inverts, so misrepresents, the order of justification.

What about cases when Q is metaphysically heavyweight? Here Pritchard takes an interestingly different tack. It is notable that there is again little if any prospect of construing the evidential bases in the columns for the *Other Minds*, *The Past* and *Induction* inferences schematised in the above table as factive.[30] And amplification of background evidence—even prescinding from the misgiving just voiced about the relevance of the tactic in general in defence of dynamic closure—looks to be forestalled as well. When the issue is whether there are other minds at all, or whether the remote past was a reality, or whether the world continues, as in the past, to be inductively amenable, what background evidence might we possibly adduce?

Pritchard's ingenious response is to query whether, when it comes to "Hinges" of so fundamental and general a character, it is proper to think of our attitude of unconditional certainty as a species of *belief*. Dynamic closure, as he formulates it, talks about the knowledgeable status of a *belief* formed by competent deduction from a known proposition. If it is actually impossible to *believe*, properly so described, the conclusion of a competent deduction of a metaphysical heavyweight from a known non-metaphysical heavyweight, then the antecedent to Dynamic Closure, appropriately formulated, is, necessarily, never fulfilled in such a case. So the apparent counterexamples to the principle are finessed.

What to make of this move? I have some sympathy for the idea that, as it features in ordinary thought, the notion of belief is somewhat generic,

indifferent to the distinctions between religious faith, unevidenced trust, the kind of attitude that is generally reckoned to be an essential ingredient of knowledge, and—maybe the same thing—the kind of attitude that one may cautiously profess in lieu of claiming knowledge. In earlier work I proposed that our attitude to the heavyweight hinges should be regarded as one of unconditional but rational *trust*. That we *know* such propositions to be true is arguably an inappropriate idea if they indeed are out of reach of any probative cognitive achievement, lying "apart from the route travelled by enquiry", as Wittgenstein put it.[31] In a similar vein, Pritchard proposes that we sharply distinguish the kind of attitudinal acceptance that is potentially a component of knowledge—what he terms *K-belief*—from other kinds of acceptance and contends that it is only and exactly K-belief for which Dynamic Closure properly holds.

That said, I would suggest that the attempt to protect Dynamic Closure in this way is vulnerable to a charge of creative accounting. At least it is so unless or until we have some independent account of the contrast between the two kinds of attitude—K-belief and the kind of acceptance we go in for with metaphysical Hinges— from which their differential behaviour with regard to Closure follows and is thereby independently explained. The fact is that a theorist who rejects Dynamic Closure on the kind of grounds rehearsed in the preceding is claiming that one cannot always earn a rational epistemic warrant for accepting a proposition by deducing it from something which one knows. It is no contradiction of the point of this claim to reply that, to the contrary, one can indeed always so acquire warrant for *believing* such a proposition, provided it is the kind of proposition which admits of belief in the first place. Absent some independent account of the relevant notion of belief and a justification for partitioning off the exceptions as non-believable, the claim that Dynamic Closure holds comes to little more than the claim that it holds—except when it doesn't.

## 6 Conclusion

It is not a good reason to maintain Dynamic Closure that we otherwise put in jeopardy our ability to advance our knowledge and justified beliefs by deduction. There is no such jeopardy provided we can corral and theorise the exceptions. The hypothesis I would propose is that they are one and all cases where the conclusion of a valid piece of reasoning encodes an authenticity condition for the warrant we take ourselves to have for one or more of the premises. When that is so, the reasoning ought rationally to be regarded as powerless to induce justified confidence in its stated conclusion since absent such confidence in advance, one rationally ought to lack confidence in the putative warrant for the premises. That, in a nutshell, is how and why transmission failure occurs,

Cases where the premises are known are no exception to this possibility. Dynamic Closure, in full generality, is false provided knowledge of the conclusion of a valid inference from known premises is required to be *rationally* based thereon.

Better learn to live with it.

## 7 Coda—Static Closure

What of static closure principles, in particular *Knowability Closure*? This was formulated as follows:

If P is knowable and may be known to entail Q, then Q also is knowable.

It should be apparent that even this principle fails. One group of counterexamples will be cases where Q is both entailed by P and is a prop for the authenticity of the means—evidence or cognitive process—whereby P may be known but which lies beyond all independent corroboration—the situation of the metaphysical heavyweights. Another theoretically possible type of counterexample will be where the authenticity condition, Q, again enjoys the status of a prop for the acquisition, by the relevant method in context, of knowledge that P but allows only of independent corroboration falling short of the strength required for knowledge.

Dretske and Nozick did not orchestrate their respective discussions of knowledge closure around the distinction between static and dynamic versions. Nor, unless one believes that it is indeed constitutive of knowledge that it be counterfactually *sensitive* in their technical sense, did they offer cogent reason for rejecting either version of the principle. I conclude nevertheless that their instinct was sound. So long as knowledge by deduction is subject to a rational basis requirement, Dynamic Closure fails in full generality. And so long as knowledge generally is required to be the product of a specific cognitive achievement, Static Closure founders on the status of the Heavyweight Hinges.[32]

## Notes

1 This volume, p. 8
2 This volume, p. 114. I'll focus on Jope's formulation in what follows.
3 Pritchard (2016a), p. 14.
4 Page reference this volume, p. 8.
5 Page reference this volume, p. 8.
6 Wright (2000) onwards — see eight further entries in References.
7 John Pollock says "undercutting". *Locus classicus*: Pollock (1974).
8 For ease of exposition, I will sometimes omit explicit reference to the "faculty" case in the formulations to follow.
9 I conjecture, though I have no space to investigate the matter fully here, that this proposal has the resources to explain the intuitively convergent ideas about transmission failure put forward in their respective contributions by Yuval Avnur and Krista Lawlor in the present volume (Avnur, pp. 86–88; Lawlor, pp. 30–33). Call any authenticity condition that for

a certain range of cognitive projects is normally considered as something which may permissibly be taken for granted without special investigation (if indeed investigation would be so much as possible), a *prop*. Then I think Avnur's idea of certain kind of epistemic method's being "biased towards" the obtaining of a certain condition comes down to our treatment of that condition as a prop for investigations conducted by that method. Likewise, in consideration of Lawlor's proposal in terms of relevant and irrelevant alternatives, I conjecture that in a wide range of cases, for a proposition Q, inconsistent with P, to be an irrelevant alternative to P for the purposes of an investigation whether P, is for its negation, ~Q, to be included among the props for the method of investigation concerned or the significance of the evidence it discloses.

10 And the same applies, naturally, to Pritchard's Competent Deduction Closure.

11 Moore (1939).

12 McKinsey (1991).

13 Putnam (1981), chapter 1.

14 Dretske (1970, 2005).

15 Pryor (2000).

16 I have responded to Pryor in Wright (2007) and (2012).

17 McDowell (1994), p. 26.

18 I am aware of course that Disjunctivists characteristically think they have the resources to address this issue. Their response is, roughly, that although BAD experiences can in principle be subjectively indistinguishable from normal perceptual experiences, an agent who is perceiving external material things in a normal way can, just on that account, know that they are, so that the sceptical argument transitions illicitly from "An agent in BAD need not be able to tell that they are" to "An agent in GOOD need not be able to tell that they are". Compare, "If drunk, I need not be able to tell whether I am unfit to drive", and "If sober, I will be able to tell whether I am unfit to drive".

There is strain in the comparison. Being drunk is phenomenologically, I am assured, nothing like normal sober consciousness. But being in one of those BAD states that is conceded to be phenomenologically indistinguishable from a possible episode of normal perceptual experience is a fortiori *very much* like such a normal episode! Yet although, while in the former, one would not be able to tell that one was, yet if in the latter, it would be clear that that was the situation. How?

One way or another, what Disjunctivism needs to claim is that E in *Ante-Moore* is tendentious — that E will be false if I am in GOOD: for, the claim must be, it can be and normally is luminous to a normally perceptually functioning subject in GOOD that such is their situation.

I have difficulty distinguishing that assertion from a Johnsonian *argumentum ad lapidum*.

19 Boghossian (1997).

20 Brain-in-a-vathood *a la* Putnam is thus intended as a metaphor for the massive error and ignorance which metaphysical realism entails, on Putnam's account, is metaphysically possibly our actual predicament. If it can be shown that, necessarily, we are not brain-in-a-vat, it is thereby intended to follow that we are, necessarily, in no such predicament, so that metaphysical realism is, necessarily, a misconception of the relationship between our thought and the real world.

21 A recent further discussion of the issues is Thorpe and Wright (forthcoming).

22 —— And in intuitionist logic too, if we take it that F and G are decidable.

23 See McGlynn (2014), section 4 for his own expression of it.

24 —— nudged by Sven Rosenkranz.
25 Britain's famous speaking clock was introduced in 1936 and still keeps perfect time for the 30 million people who use the service each year.
26 Namely, that the world did *not* spring into existence five minutes ago, albeit replete with apparent traces of a much more ancient history (Russell, 1921, p. 159).
27 Pritchard (2016a, 2016b) and this volume.
28 "Finally … I want to highlight that the notion of knowledge that is in play here is rationally grounded knowledge. Depending on one's wider epistemology, rationally grounded knowledge might well be the only kind of knowledge that there is", this volume, p. 114.
29 See note 18 above.
30 *Pace* John McDowell who once wrote that.

> We should not jib at, or interpret away, the commonsense thought that, on those occasions that are paradigmatically suitable for training in the assertoric use of the relevant part of language, one can literally perceive in another person's facial expression or his behaviour, that he is (for instance) in pain, and not just infer that he is in pain from what one perceives.
>
> (McDowell, 1978, p. 135)

31 *On Certainty* §88 (Wittgenstein, 1969).
32 Thanks to Sven Rosenkranz and to colleagues in the Stirling Work-in-Progress seminar for helpful comments and discussion, and to Matt Jope and Duncan Pritchard for giving me the opportunity to contribute to the present volume.

## References

Yuval Avnur (This volume) "The Skeptical Paradox and the Generality of Closure (and Other Principles)".

Paul Boghossian (1997) What the Externalist Can Know a Priori", *Proceedings of the Aristotelian Society* 97 (2) pp. 161–175.

Fred Dretske (1970) "Epistemic Operators", *Journal of Philosophy* 67 (24) pp. 1007–1023.

Fred Dretske (2005) "Is Knowledge Closed under Known Entailment? The Case against Closure" in Matthias Steup & Ernest Sosa (eds.), *Contemporary Debates in Epistemology* Oxford, UK: Blackwell, pp. 13–26.

John Hawthorne (2005) "The Case for Closure" in Matthias Steup & Ernest Sosa (eds.), *Contemporary Debates in Epistemology* Oxford, UK: Blackwell, pp. 26–43.

Matthew Jope (This volume) "Intuitive Closure, Transmission Failure and Doxastic Justification".

Krista Lawlor (This volume) "Modest Foundations for Perceptual Knowledge".

Stephen Luper (2016) 'Epistemic Closure' in Zalta (ed.), *The Stanford Encyclopedia of Philosophy*. URL = https://plato.stanford.edu/archives/win2018/entries/closure-epistemic/.

John McDowell (1978) "On 'The Reality of the Past'" in C. Hookway & P. Pettit (eds.), *Action and Interpretation* Cambridge: Cambridge University Press, pp. 127–144.

——— (1994) *Mind and World* Oxford: Oxford University Press.

Aidan McGlynn (2014) "On Epistemic Alchemy" in D. Dodd & E. Zardini (eds.), *Scepticism and Perceptual Justification* Oxford: Oxford University Press, pp. 173–189.

Michael McKinsey (1991) "Anti-Individualism and Privileged Access", *Analysis* 51 (1) pp. 9–16.
G.E. Moore (1939) "Proof of an External World", *Proceedings of the British Academy* 25 (5) pp. 273–300.
Robert Nozick (1981) *Philosophical Explanations* Cambridge, MA: Harvard University Press.
John Pollock (1974) *Knowledge and Justification* Princeton, NJ: Princeton University Press.
Duncan Pritchard (2016a) *Epistemic Angst* Princeton, NJ: Princeton University Press.
——— (2016b) "Epistemic Angst", *Philosophical and Phenomenological Research* 93 (3) pp. 70–90.
——— (This volume) "In Defence Of Closure".
James Pryor (2000) "The Skeptic and the Dogmatist" *Noûs* 34 (4) pp. 517–549.
Hilary Putnam (1981) *Reason, Truth and History* Cambridge: Cambridge University Press.
Bertrand Russell (1921) *The Analysis of Mind* London: G. Allen & Unwin.
Joshua R. Thorpe & Crispin Wright (Forthcoming) "Putnam's Proof Revisited" in Sanjit Chakraborty & James Conant (eds.), *Engaging Putnam* Abingdon: Routledge, pp 63–88.
Timothy Williamson (2000) *Knowledge and Its Limits* Oxford: Oxford University Press.
Ludwig Wittgenstein (1969) *On Certainty* (edited by G.E.M. Anscombe & G.H. von Wright) New York and London: Harper Torchbooks.
Crispin Wright (2000) "Cogency and Question-Begging: Some Reflections on McKinsey's Paradox and Putnam's Proof", *Philosophical Issues* 10 pp. 140–163.
——— (2002) "(Anti-)sceptics Simple and Subtle: G.E. Moore and John McDowell", *Philosophy and Phenomenological Research* 65 (2) pp. 330–348.
——— (2003) "Some Reflections on the Acquisition of Warrant by Inference" in S. Nuccetelli (ed.), *New Essays on Semantic Externalism and Self-Knowledge* Bradford: MIT, pp. 57–77.
——— (2004) "Warrant for Nothing (and Foundations for Free)?", *Proceedings of the Aristotelian Society, Supplementary Volume* 78 pp. 167–245.
——— (2007) "The Perils of Dogmatism" in Susana Nuccetelli & Gary Seay (eds.), *Themes from G. E. Moore: New Essays in Epistemology and Ethics* Oxford: The Clarendon Press, pp. 25–48.
——— (2011) "McKinsey One More Time" in Anthony Hatzimoysis (ed.), *Self-Knowledge* Oxford: Oxford University Press, pp. 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: Oxford University Press, pp. 451–486.
——— (2014) "On Epistemic Entitlement II: Welfare State Epistemology" in D. Dodd & E. Zardini (eds.), *Scepticism and Perceptual Justification* Oxford: Oxford University Press, pp. 173–189.