sciendo

The Information-Theoretic Account of Knowledge, Closure and the KK Thesis

James Mattingly
Georgetown University

doi: 10.2478/disp-2022-0007 BIBLID: [0873-626X (2022) 65; pp. 105–32]

Abstract
One common objection to Dretske's Information Theoretic Account of Knowledge (ITAK) is that it violates closure. I show that it does not, and that extant arguments attempting to establish that it does rely instead on the KK thesis. That thesis does fail for ITAK. I show moreover that an interesting consequence of ITAK obeying the closure principle after all is that on this view if skepticism is false, we can have a great deal of empirical knowledge, but it is in principle impossible to know that skepticism is false. In short, a proper understanding of how ITAK closes off the KK thesis shows that we can 1) take seriously the skeptic, we can 2) respond to her appropriately that we do have knowledge and we can 3) keep closure.

Keywords
closure; Dretske; KK thesis; information; skepticism

1 Introduction

I will start with a confession: I love Dretske's Information Theoretic Account of Knowledge (ITAK) and I think you should too. It is elegant, intuitively obvious and it solves many of the traditional problems of epistemology naturally and easily. Now I am under no illusions that all, or even a large fraction of, philosophers agree first that what separates knowledge from merely true belief can be cashed out in causal terms, and second that ITAK is a good choice if so. While I cannot here hope to convince those with broadly internalist commitments to

make the switch, I *will* try to have something useful to say even to those with intuitions disjoint from mine. In particular my analysis will show that those who take seriously the challenge of skepticism, who think we do have empirical knowledge and are committed to unrestricted logical closure can find an epistemological home in ITAK.

Based on informal (read: conversational) investigation it seems that many who would otherwise see significant advantages to ITAK reject it because it appears to violate the closure principle. I am not sure that I am immovably committed to closure myself, but I do regard violating closure as prima facie grounds to reject an epistemological theory. Interestingly, and despite a strong consensus to the contrary, ITAK does not violate that principle. This has been (sort of) shown by Jäger and Shackel and (sort of) conceded by Dretske. However the participants in their exchange (and everyone else) seem to have concluded that it does violate closure after all, or rather should be made to. I will rehearse that exchange here, and offer a streamlined and more compelling argument showing that ITAK obeys the closure principle. To those who reject ITAK only because of their commitment to closure, I say, "welcome back."

The key philosophical innovation here, however, is that by clarifying exactly how ITAK obeys closure, and concomitantly how it violates instead the KK thesis,[1] one can discover yet another beautiful fact about the theory: if skepticism is false, we can correctly *assert* it to be false and yet we cannot, *in principle*, correctly assert that we know that it is false. This is a significant result and it shows that the theory situates its adherents in a uniquely comfortable position: the skeptic need not be refuted, that is to say have her worries ruled out of court by the right theory of knowledge (as is, I believe, impossible in any case), but can be comfortably responded to without question begging. Thus the goods that Dretske thought could only be had by abandoning closure, are readily available with it. Those who, like Dretske, only reluctantly abandoned closure because doing so seemed a necessary concession in finding a middle ground between accepting skepticism (intolerable) and ruling it out altogether from the outset (impossible), should find this a welcome result. Others should welcome the result as well just because it is an intrinsically interesting feature of the theory.

Now, an important, but by no means the sole, distinction between internalist and

---

[1] For my purposes this is the thesis that if one knows that-p then one is always in a position to come to know that one knows that p. This is a comparatively weak version of the thesis and so the fact that it is false on ITAK is a strong result.

externalist theories of knowledge is that, generally, internalist views take the responsibility of agents as relevant to their knowledge state while externalist views do not. This responsibility condition is important for the internalist in separating those who know from those who merely believe correctly.[2] Externalists by and large attempt to separate these groups by a condition of truth-tracking, or reliability, or something of that nature, appealing to some feature of the way the agent is embedded in a knowledge-apt situation. This is all well-known, but should be kept in mind as we evaluate ITAK, for we need to be clear on what the theory itself entails before evaluating whether we think we want to adopt the theory given that it entails those things.

1.1 Closure and the specter of skepticism

To those who are moved but not persuaded by the skeptics' arguments it is deeply troubling that we might either have to accept skepticism or come to rule it out of court altogether. Those who are not skeptics but also do not reject skepticism out of hand insist that there must be some middle ground where we can accept that global skeptical claims are irrefutable, while at the same time enjoying our everyday, as well as scientific, empirical knowledge. These thinkers take it that whatever epistemological theory we are considering, if it makes us so strongly convinced that we have hands that we rule out the possibility that we are brains in vats, or contrariwise makes our conviction that we *cannot* rule out that we are brains in vats incompatible with our knowledge that we have hands, then that theory should be abandoned. Such thinkers (Nozick and Dretske prominent among them) take us to have what amounts to *equally* strong assurance that we have hands *and* that we cannot rule out the possibility that we are brains in vats (or some other global skeptical hypothesis). But, they take it, a natural solution presents itself: accept both premises (we know we have hands; we cannot rule out global skepticism) and reject the inference that brings them into conflict. This solution is at the heart of Nozick's truth tracking theory, Dretske's conclusive reasons view and (apparently) Dretske's Information Theoretic Account of Knowledge (ITAK). The strategy is to develop a theory of knowledge on which there is no proper inference from "I know I have hands" to "I know I'm not a brain in a vat" despite first knowing that "if I have hands I'm not a brain in a vat" and then inferring the latter claim on that basis. In other words, the solution is to develop a theory of knowledge that violates closure.[3]

---

[2] See Greco (2010: ch. 3) for a different account of what separates internalists from externalists.

[3] Or at least make closure failure a necessary condition on a good theory.

For my purposes, closure will be the principle that knowledge is closed for beliefs formed (or sustained) under appropriate operations of logical entailment. In particular suppose that an agent knows that p, and knows that $p \rightarrow q$. Then that agent on drawing the conclusion q, based on both p and $p \rightarrow q$, and either coming to believe q on that basis or having the belief q sustained on that basis, knows that q. Similar considerations hold for other logical operations. Because closure appears to connect ineluctably with how we understand the logical structure of our store of knowledge, resistance to theories that reject it can be fierce. Dretske outlines some of these in his (2005) where expressions like "abominable", "least plausible" and "bizzare (sic)" abound, and where the fact that one's view entails it provides a "devastating objection" or a "*reductio ad absurdum*" of the view. (17) It is easy to be sympathetic to the thought that Moore type paradoxes like "I know this is an acid but I do not know that it is not a non-acid" are both intolerable and consequences of any view that rejects closure.

I will not here rehearse the many arguments for and against epistemic closure. Let me merely point out that even its strongest critics oppose it for extrinsic reasons. Dretske for example regards it as an obvious principle that ought to be accepted by any rational person *unless* very strong reasons make us abandon it. For example he says, "closure sounds like an eminently plausible principle. Everything else being equal, then, we ought to keep it. But everything else isn't equal" (Dretske 2005: 18) because denying it is, he thinks, the only way to avoid skepticism. If being left with nothing to say to the skeptic is really the best reason for rejecting closure, then my conclusions should be very welcome because they show how accepting closure does not leave us tongue-tied before the skeptic. Happily there are no similarly strong reasons to accept the KK thesis, and the reasons for rejecting it here are integral to the theory's basic story about how knowledge works. We will get to all of that shortly, but now let us turn to the basic structure of ITAK.

2 The information theoretic account of knowledge

ITAK has two defining features as laid out by Dretske in his *Knowledge and the Flow of Information*. The first is its account, in terms of information flow, of when an agent, K, knows; the second is an account of when a signal causes information to flow to an agent.

> *Knowledge*: K knows that s is F = K's belief that s is F is caused (or causally sustained) by the information that s is F. (Dretske 1981: 86)

> *Informational Content*: A signal *r* carries the information that *s* is F = The conditional probability of *s* being F, given *r* (and *k*) is 1 (but, given *k* alone, less than 1). (Dretske 1981: 65)

(Here "*k*" stands for K's background knowledge.) As Dretske points out, it is not easy to give a precise characterization of what it is for information to cause a belief. I will not take up that challenge here. For current purposes that should not much matter.[4,5]

In addition to these two defining features it is necessary to say something about the circumstances in which such signals can move conditional probabilities from less than 1 to 1. All potentially information bearing signals travel in channels. These can be composed of wires, light signals, sound waves, boxes in delivery trucks, the memory and running legs of Pheidippides, etc. In order that such signals *do* bear information many things must be in order: wires cannot be broken; there must be no opaque barriers; etc. When everything is in place to sustain information bearing signals, the channel is called a *communication channel*.

> The channel of communication = that set of existing conditions (on which the signal depends) that either (1) generate no (relevant) information, or (2) generate only redundant information (from the point of view of the receiver). (Dretske 1981: 115)

Example: an electric cable in a house might be part of an information channel when it connects the pressing of the doorbell button with the ringing of the bell. That same cable might be what I am getting information *about* when I am testing the broken doorbell and seeing whether current flows in that part of the system. In the former case the continuity of the cable is a condition on which the signal depends, and that continuity gives no relevant information; in the latter case the continuity (or discontinuity) of the cable is what we are getting information about. As we will see, it can be tricky to evaluate this distinction in practice.

---

[4] A common sticking point here is the difference between an agent's *assessment* of the probability, and the probability itself. The conditionals here refer to the latter: the question is whether when *this signal* in *this channel* is present the probability of s being F is 1 *not* whether it is ever a good idea for an epistemic agent to appraise the probability as 1 (as it never is for any empirical knowledge claim).

[5] Another common sticking point is that this view seems to have no room for knowledge of tautologies, because those have probability 1 on any background knowledge set. The standard reply is that this is a theory of *empirical* knowledge. That is probably sufficient. My own (controversial) view is that tautologies have no content and thus are not objects of knowledge. In what follows I will be considering only contingent propositions.

3 Calibration and communication channels[6]

With those three things in place we are prepared to evaluate the theory with respect first to closure, then the KK thesis and finally with respect to its possible engagement with the skeptic. However, before confronting these matters head on I will do a little preparatory ground clearing to say more precisely what is necessary for properly functioning channels.

"Calibrate", like "being", is said in many ways. For our purposes, though, the salient ways are two: to calibrate an apparatus is to make it suitable as a probe of some feature of the world. For example, to calibrate a thermometer is to prepare it so that its readings reflect the temperature of its surroundings. That is Sense 1. Sense 2 is different, and it takes it that to calibrate an apparatus is to *check* that it is suitable as a probe of some feature of the world. These senses are by no means independent of each other, but they are distinct.[7] However they are easily conflated with each other. This is especially true because generally our method of calibrating in Sense 1 is to tinker around with the apparatus until calibration in Sense 2 is successful.

Calibration applies more generally, not only to standard measurement tools. In fact it applies to anything that counts as a channel of communication, and the distinction is the same as that between a channel being suited to carry information and some agent knowing that it is so suited. Arguments concluding that ITAK in particular and tracking theories generally,[8] should fail to obey closure are, as we shall see, predicated on conflating these senses of calibration.

Consider the famous zebra/painted mule example. It is always possible that we are at a zoo whose keepers want to keep the zebra pen open for view on zebra bath day and so they tend on bath day to paint mules to look like zebras and let them loose in the zebra pen.[9]

---

[6] I elaborate on all of this in much greater detail in Mattingly (2021).

[7] They are distinct in roughly the way that externalist views of epistemology are distinct from internalist views.

[8] These are theories on which the justification condition in the justified true belief conception of knowledge is replaced with a requirement that belief formation is properly responsive to the way the world is. So, one might count beliefs as knowledge if they are true and would not be believed were they not true.

[9] In either the sense that possibly today is the day that this zoo with this practice has washing day, or the sense that possibly we are at such a zoo because such zoos exist even though in fact this is not one such. Plausibly many (not me) will only have the intuition that knowledge is possible in the latter case. Thanks to John Greco for urging clarity on this point.

The difficulty for a truth-tracker is to accommodate our intuition that one can be at such a zoo, see zebras and know they are zebras. Why is this a difficulty? Because of the possibility that, for example, those animals are painted mules. Since apparently we cannot know that they *are not* painted mules, then allowing closure and applying modus tollens would let us conclude that we *do not* know that these are zebras.

But closure is not the problem here. Rather the problem is a confusion between the two senses of calibration. Naturally to establish that these animals are not painted mules, and so to eliminate from consideration that possibility, would require knowing that the signal that these are zebras is carried by a properly calibrated information channel. But one need not know this in order for it to be *true* that the channel is properly calibrated. It is this latter fact that guarantees both that the signal carries the information that they are zebras *and* that they are not painted mules. Either the channel can carry the information or it cannot. If so we can know by its means, if not, not. Its being calibrated, however, does not let us know that it is so: Sense 1 does not entail Sense 2. To find that out we have to find out, for example, whether this is a zoo that paints mule (and perhaps whether it is bath day), maybe by interviewing the staff, or by looking at a schedule, or by some other means. But as we will see all alleged failures of closure are best explained this way: information did not really flow because the probability = 1 constraint on information flow was not met. What advocates for closure violation do not recognize is that when that condition is not met it not only falsifies claims derived by the application of closure, but *also* the original claim from which they derive. Put differently: we no more appeal to the proper working of the informational connection between zebra appearances and zebras when we come to believe they are non-mules than we do when we come to believe they are zebras.

Here then is a suggestion that should be embraced by all those who require (i) that our theory allow that we have empirical knowledge and (ii) that our theory not build in the failure of skeptical possibilities. All we need here is to keep in mind that we can know many things without knowing that we know them: the former requires only that channels carry information (calibration in Sense 1); the latter is *also* to know that the channels carrying such information do so (calibration in Sense 2).

It is tricky, but critical, to keep the following claims distinct: (1) p entails q; (2) *the fact that I know that p* entails q; (3) *the fact that I know that p* entails *that I (am in a position to) know that q*. Dretske especially has been quite clear on why the fact that I know that p cannot be used to generate every entailment of the fact that I know that p. In part, at least in the case of ITAK, this is because even though the fact that I know that p does entail the fact that appropriate channel conditions supporting the flow of information that p were in

place, this entailed fact is not necessarily contained in any signal I have received—and it is *necessarily* not contained in the signal that bore the information that p. He and others are quite right that the fact that I know that p coupled with the fact that my knowing that p itself entails q does not suffice to give me the knowledge that q. But where these thinkers have gone astray is in thinking that such restrictions are relevant to the closure principle. They are instead relevant to the fact that knowing does not entail knowing that you know. As we will see the KK thesis *does* fail for ITAK, as it should.

In the case of zebras at the zoo, I *could* try to reason this way: this is one of those occasions on which visual impressions *like* those of looking at a zebra would bear the information that a zebra is before me. Thus the visual channel between me and that creature over there is information bearing. And that rules out masquerades, including mules painted like zebras. So that creature is not one of those. Here I would be reasoning first from my knowledge that my channel is information bearing, and that would indeed be an illegitimate inferential strategy (unless I had independently extra knowledge about the information channels at the zoo). But that is not an application of closure, and is not how one does reason when appealing to closure.[10] Instead closure works like this: one believes it is a zebra, knows that zebras are not painted mules and comes to believe it is not a painted mule on the basis of those two facts. If the channel is working then one knows it is not a painted mule (because in fact one has the information); if the channel isn't working, one does not know it is not a painted mule, but nor does one know it is zebra in that case either (because one does not have the information).

Consider an example from Dretske:

> Tree rings carry information about the age of the tree. Suppose the tree stump has twenty rings. Assuming everything normal, this carries the information that the tree is 20 years

[10] I have seen the following argument: "I know that my tank is full (by looking at the gauge) and I know that the gauge (my channel of information) indicates that my tank is full. Can I thus come to know by simple inference that my gauge is working correctly (and after repeated exercise come to know, again by simple inference, that my gauge is reliable)? It seems not." Correct, but closure does not say it should. If I do not already know *that* I know my tank is full, then I do not know that my basis for that belief (the gauge indicates it) bears the appropriate information. What I know, at the moment, is that the tank is full not further things about the channel bringing me the information. So I do *not* know that I know it. This is a simple distinction but one that is apparently very tricky to keep in mind. A way to know that the gauge is reliable is by seeing that it indicates full, checking that the tank is full, repeating and then drawing the inductive conclusion that it is reliable. That however is *also* not an application of closure.

> old. That the tree is 20 years old implies, logically, that the past (at least 20 years of it) is real. So Russell's skeptical hypothesis about the world being only a few minutes old must (logically) be false. Do the tree rings carry the information that Russell's hypothesis is false. No! Why not? Because information, as we ordinarily understand it, and as my own theory of information was designed to show, is not closed under logical entailment. (Dretske 2006: 410)

The conclusion that we cannot learn that the past is real *from tree rings* is correct, but we can learn it *from the fact that tree rings bear information about tree age* in conjunction with those rings. If I know that I know that the tree is 20 years old, then I can use that to find out that the past is real. I do not get that information now, but I can come to believe this fact about the past by thinking about tree rings I suppose. One normally does not proceed that way, but lots of things can make me come to believe something that is then causally sustained by the information that it is the case.

That one *could* come to believe something by appeal to the working of a channel does not entail that one *does* do so. More importantly, appealing to the well-functioning of a channel is simply not an application of closure, and the fact that such appeals are sometimes illegitimate does not ever make appeals to closure illegitimate. The kind of appeal to the working of a channel we saw above is implicitly appealing to the KK thesis, and that has nothing to do with closure.

The idea that channel conditions restrict the information that is available for signaling is tempting, but is not quite right. What is right is, as we said above, that not *everything* that is entailed by the fact that some signal bears the information that p is available for signaling by the channel that the signal is in. In particular, the fact that the signal bears the information, and the fact that the channel is apt for signal bearing, and related facts are not available for signaling. However anything entailed alone by the fact that p *is* available for signaling, and will be available for signaling to the recipient provided she does not already know it, and provided that she does know that p entails it.[11]

---

[11] More generally anything only entailed by p in conjunction with other things is available for signaling as long as those things are not themselves conditions on the channel that bears the information that p.

4 ITAK does not violate the closure principle

Now we can move beyond those primers of intuition, and directly consider the question of whether ITAK itself violates closure. In this section I will first argue directly that it does not. I then consider Dretske's arguments that it does and reject them. Finally I consider attempts by Adams and his collaborators to show that closure fails and explain why instead it is their arguments that fail.

4.1 Directly

Suppose I know that p, and I know that $p \rightarrow q$, and I draw by way of modus ponens the inference that q and come to believe that q on that basis. Let us apply ITAK and see whether I am positioned to know that q.

If my background knowledge does not entail q, then the issue is easily settled. Prior to receiving the signal that p, my background knowledge alone gives the probability that q as less than 1, but my background knowledge alone gives the probability that $p \rightarrow q$ as 1. The signal that p raises the conditional probability that p to 1. The laws of probability then entail that the probability that q is 1 as well. Thus the signal that p bears the information that q. Since I come to believe that q based on the information that q, I know that q.

On the other hand, suppose that my background knowledge does entail q. Then if I were to reflect on what I already know, I could come to know q. Why? If all of my empirical knowledge is gained via information bearing signals as ITAK suggests, then at some state my background knowledge *did not* entail q, else I could not have received the information. When I acquired the information that gave me the knowledge that does entail q, then that entailment would support an inference (and my subsequent coming to believe) that q from the signal that bore that information, because it also bore the information that q. Here is why. Before receiving whatever signal it may have been, my background knowledge, k, did not make the probability that $q = 1$. But suppose that in light of a new item of knowledge, call it p, together with my background knowledge, did make the probability that $q = 1$. Then, $(k \,\&\, p) \rightarrow q$. And so any signal that bears the information that p also bears the information that q. An optimal epistemic agent could then know that $(k \,\&\, p) \rightarrow q$, and could therefore infer that q on its basis. Whether or not the agent would know that the inference is on the basis of the information that q is irrelevant, it is still the case that it *is* on that basis. Thus, were I to have exercised fully my epistemic capacities, I could have inferred (and come to believe) that q on the basis of the information that q. Therefore I would be positioned to know that q.

So on ITAK if I know that p and I know that $p \rightarrow q$ then I am positioned to know that q. And moreover the only reason I would fail to know that q is because I in fact failed to draw a conclusion, and *not* because ITAK violates the closure principle. It does not.

4.2 Contra Dretske

A typical argument that closure fails goes like this:

> If knowledge that P requires one (or one's evidence) to exclude not all, but only all relevant, alternatives to P, then, it seems, one is committed to a failure of closure. The evidence that (by excluding all relevant alternatives) enables me to know there are cookies in the jar does not enable me to know that they (what I see in the jar) are not papier mâché fakes since papier mâché fakes are not (usually) a relevant alternative. So the evidence that gives me knowledge of P (there are cookies in the jar) can exist without evidence for knowing Q (they are not fake) despite my knowing that P implies Q. So closure fails. (Dretske 2005: 19)

However this is not quite the same as an argument that closure fails for ITAK itself. For while a "relevant alternatives" view of the standards for knowing is part of Dretske's explanation in *Knowledge and the Flow of Information* for why we should adopt ITAK, it is simply not part of the formal theory itself. What ITAK offers instead is an account of information channels and the conditions under which they suffice to carry information, the so-called channel conditions. The question here is whether channel conditions function in the same way as relevant alternatives in the above argument.

An exchange in *Erkenntnis* in 2006 comes very close to settling this issue. There Baumann (2006) presents the only published calculation that I am aware of that purports to show that ITAK is not closed. Baumann's idea is that information flow itself is not closed on Dretske's definition of information bearing signals. Indeed he is quite right about this. For if a signal bears me the information that p and I know that $p \rightarrow q$, the signal may not bear the information that q. How? Well as we saw above, if I already know that q or my background knowledge itself entails q, the signal does not convey that information to me. Thus information *flow* is not closed. Baumann's mistake is to neglect to note that this failure of information flow *only* happens in exactly those cases where my background knowledge already entails q. My argument above shows that even in those cases, I am, after all, in a position to know that q, and so while information *flow* is not closed, information itself

is and,[12] crucially, so is knowledge. Put another way, whether information flow is closed or not, if I have the information that p and the information that p → q, then either the information that q flowed in at some earlier point in my cognitive history, or it flowed in along with the information that p.[13]

Shackel (2006) realizes essentially this and analyzes Baumann's error. He comes close to convincing Dretske that ITAK really is closed. Shackel proposes an Information Closure Principle that he takes to apply to ITAK.

> *Principle of Informational Closure*: If a signal, S, gives you information that E, you know that E entails G, and G is available for signaling, then S gives you information that G. (Shackel 2006: 396)

The guiding idea is simply that if you already know something then no signal can inform you of it, and that is just an application of Dretske's condition on signals bearing information. Both Dretske and Shackel then take this principle to advance the debate in different directions however. Shackel supposes that he can show that any one piece of empirical knowledge suffices to establish that skepticism is false. This is because he thinks the proposition that skepticism is false is available for signaling. As we saw above Dretske would find this conclusion unwelcome. His response is to accept something like the Information Closure Principle and to deny that channel conditions are available for signaling by appeal to a new restriction on how signals bear information:

> Given background knowledge, a signal can raise the probability of Q to one and still not carry the information that Q. When Q is a condition of the channel on which the flow of information about P depends – when, that is, Q is or involves the existence of what I called a channel condition—this is not true even if (by an appropriate description of Q) P implies Q. (Dretske 2006: 413)

Dretske is correct that the fact that some channel is properly functioning is not available for

---

[12] This is because if I have the information that p and the information that p → q then if I correctly draw the inference that q either I already have the information that q or drawing the inference bears that information.

[13] A helpful referee points out that this distinction is very close to Crispin Wright's between closure and warrant transmission in, for example, "Some reflections of the acquisition of warrant by inference" (Wright: 2003).

signaling, but we do not need this new restriction to prevent that. His original condition on information flow suffices.

Recall the notion of the communication channel. It is part of the definition of the channel that it comprises all of those conditions that make it possible for signals to carry information. In order, now, to see how the original view closes off signals from bearing information about the channel conditions that support them, let us imagine receiving some signal. Does that signal bear information about its source? Suppose it does.[14] Then the conditional probability in the *absence* of the signal that the source is in the state it is in must be less than 1, and it must be 1 in its presence. This tells us right away that the channel conditions are met. But if they are met then it must be the case that *prior to transmission of the signal* the probability that, were the signal to be transmitted it would raise the probability at the receiver to 1, is already 1. For suppose it were less than 1. Then were the signal to be transmitted to the receiver, it would not raise the probability to 1 on pain of rejecting our supposition. This is not to say that no signal that bears some information about a source can bear *any* information about the channel in which it is traveling. It is just that the signal that bears the information *that s is F* cannot bear the information that the channel is functioning. For example the signal might, by its nature, carry the information that the channel is electric-current bearing. It may even, perhaps, bear the information that it was carried by that very channel. None of that would support the inference that the channel conditions that allow it to bear information about that source are in place from the fact that the source of the signal does indeed have the properties the signal is informing us about.

That was all pretty fast and abstract. So back to Dretske's example of cookies versus papier mâché fakes: if I look in the cookie jar I can see that there are cookies in there, and so know that there are. But Dretske thinks that I cannot know that they are not papier mâché fakes. Why? His idea seems to be that we would be inferring that the cookies are not papier mâché by appealing to the fact that the channel carrying the signal with the content that these are cookies is a good channel, and is thus carrying information. The idea would be that we know they are cookies, realize that we know it, realize that that entails that the channel is sound, and then somehow use that to infer that the cookies are not papier mâché. Agreed, we should not do that, for that would be to appeal to our knowledge of what we

[14] If not, there is nothing to worry about. The channel conditions are not in place, and clearly we cannot learn that they are from the signal.

know. But we *are not* doing that when we appeal to closure. What we are doing, instead, is inferring that the cookies are not papier mâché on the basis of the fact that they are cookies and the fact that cookies are not papier mâché. We would be wrong to conclude this, and our belief would not be based on the information that they are not papier mâché were the channel not carrying the information. But if it were not then we also would not know that they are cookies in the first place. Both the belief that they are cookies and the belief that they are not papier mâché must *rest on* the fact that the channel carries information if they are to be knowledge. But questions about channel conditions simply do not arise for us as knowers about cookies and their relation to papier mâché simulations; they arise instead for us as theorists about the conditions under which we are successfully knowing about the cookies.[15] It is a mistake and leads to much confusion to mix these two types of consideration.

I submit that Dretske should have accepted the argument about information closure *without* adding the stipulation that a signal can raise the probability of some proposition, p, to 1 without bearing the information that p as he does in his (2006). Naturally, as the creator of the theory he can modify to whatever extent he desires. But Dretske is simply wrong to assert that we do not understand information to allow learning about channel conditions, as I have argued in a variety of ways. Instead what we do not learn from a given channel is that it *is* bearing information. But Dretske's original theory already has that consequence. It is simply not true that "information, as we ordinarily understand it …isn't closed under logical entailment." (2006: 410) That is directly in conflict with the Shannon-Weaver account on which the entire discussion of information in the book is based.[16]

Moreover, despite his assumption to the contrary, ITAK is simply not a relevant alternatives theory; Dretske presumably tried to build that in, but his definitions do not make it such a theory. If anything, relevant alternatives are useful heuristics for us, as theorists, to decide whether some channel is bearing information or not. But it is the probability that is determinative of whether the world really is as the signal says it is, and that is where

---

[15] Clearly we can investigate particular channels to see whether they work. But that will always be under the assumption that empirical channels *as such* bear information. Were we in a simulation, for example, they would not and we would after all have no worldly knowledge.

[16] All we need to know here about the Shannon-Weaver account is this: information flow is essentially equivalent to a reduction in one's uncertainty about the possible ways things could be at the origin of the signal bearing the information.

the information is.

Let me make this discussion a little more clear. The new condition is not only unnecessary; it severs the precise connection to the theory of information and substitutes something vague. As Dretske hints (2006: 413) it is now a matter of under what description some claim is made that matters for whether we can come to be informed about it. I do not claim that Dretske would be happy with my analysis of his theory, nor do I claim that, on my analysis of the theory, ITAK is what Dretske was hoping to produce.[17] Here is the thing though. Dretske over the course of his career took closure violation as the only way to respond non-dogmatically to the skeptic and yet still have knowledge. It is a delicate balance to both have empirical knowledge and not to rule out skepticism from the start. As we will see however, a failure of the KK thesis is all that is necessary for that. Interestingly Dretske *almost* saw that in his (2004). It will be instructive to see what he missed, as we will in section 5.

4.3 Contra Adams et al.

Over the last few years Adams and his collaborators have done a good job clarifying many of the issues surrounding closure and truth tracking theories generally. They continue to maintain, however, that tracking theories do violate closure of necessity. They have produced a lot of fascinating work that attempts to chart precisely the scope of this violation. They are attempting to establish two interesting theses that, if true, would be very important: that the closure condition on knowledge fails generally; that failures of closure are restricted in scope. I am persuaded by neither, but will restrict my attention to the first claim, that closure fails generally for tracking theories.[18]

One of Adams' recurring examples is the use of litmus paper. Suppose I am in my chemistry lab. Then I can come to know of some liquid that it is an acid if I dip a blue

---

[17] The exchange between Dretske, Baumann, Shackel and Jäger has been noticed as well by Kipper (2016). Kipper's view accords with mine in the sense that he believes Dretske should not have added this new stipulation, but he *also* believes that Dretske's theory is fine as is, and he thinks it violates closure. On the last point he and I diverge. I do not have space to consider his fascinating proposal here in any detail. I will simply point out that it is strictly speaking a modification of Dretske's proposal, specifically the information transmission condition. Kipper's *new* theory violates closure, Dretske's does not.

[18] Naturally this is not to say that in every case knowers fail to know the logical consequences of what they know. Rather it is to say that closure cannot be appealed to on a truth tracking theory, and that there are cases where a believed and properly drawn logical consequence of what one knows is not known.

sheet of litmus paper into the liquid and it turns pink. This is true whether or not there is another possible liquid that would turn my paper pink that is not an acid. The reason is that if the channel conditions on information flowing to me about the acidic character of the liquid are in place, so that information actually can flow, then that is because the probability is zero that one of those liquids is in my lab, for whatever reason. Thus given these channel conditions, when the paper turns pink (an information bearing signal) the conditional probability of the liquid being an acid goes from less than 1 to 1. Adams et al. then wonder whether I can come to know that the liquid is *not* a non-acid that turns litmus paper pink. Their conclusion is that if so it is "surely not from the litmus paper's turning pink alone" (2012: 183). The quick reply here would be "of course not, but I also do not learn that the liquid is an acid from the paper's turning pink alone." The conditions on litmus paper being a good test for acid are exactly the same as it being a good test for not being a non-acid. That is just logic, one might think, but because of the persistent confusion around the issue, it is worth grounding this all in detail.

We are concerned with the channel that is supposed to connect pink-turning litmus paper with acidity, and what the operation of that channel tells us about closure. There are two cases to consider, one is a universe in which there are no non-acids that turn blue litmus paper pink, the other in which there are such liquids. Let us consider them in turn. (1) There is no such liquid. Then pink-turning litmus paper carries the information that it was dipped in acid. It also carries the information that it was dipped in a non-non-acid. On any tracking theory of knowledge I can come to know that I am dealing with a non-non-acid by noting that my litmus paper turned pink. No failure of closure here. The channel from *liquid is acid to litmus paper turns pink* is in good working order. (2) Such a liquid does exist. Then in order to assess whether I have knowledge that the liquid I am testing is an acid *because the litmus paper turns pink* the theorist would need to assess the probability that some of that liquid is in my lab. We cannot specify how to do this in general, but we would do it the way we assess all questions about channel conditions. While we can never be sure, perhaps, that some channel or other is really capable of transmitting information, we make our decisions about which channels support information flow and which do not by appeal to intuition, to our past experience, to our analysis of the nature of the signals involved and their means of generation, etc. All of this makes it very complicated and epistemically risky to *assert* of some channel that the context is right and that it is capable of transmitting some piece of information to some knower with some background knowledge. However if the channel does in fact support the flow of information, then that is an absolute fact about the channel, and as such the channel also supports the flow of information that is

law-like connected to the other information whose flow it supports. This is true no matter what I happen to decide using my various inductive practices. The long and short of it is that there are two sub-cases to check:[19] (A) where the channel does not function because the existence of the mystery liquid outside the lab interferes with information flow, and (B) where it does function because the mystery liquid outside the lab does not interfere with information flow. (A) Suppose the existence of the strange substance is fact interfering with the information flow in my lab. If I dip my litmus paper into some liquid or other and it turns pink, I simply do not know that it is an acid. By assumption the channel connecting *litmus paper turning pink to liquid is an acid* does not function. So despite all my training in chemistry and my good lab practices and my habit of thinking that such a channel is always in order, I just cannot come to know that it is an acid using that litmus paper. This is because in my current situation litmus paper turning pink does not carry the information that it was dipped in acid. My failure then, to know that it is a non-non-acid is not a closure failure, but is instead a failure to have appropriate knowledge to close under logical operations in the first place. (B) The existence of such a substance does not interfere with the information flow in my lab. Then the pink-turning litmus paper is at the end of a communication channel that does in fact carry the information that the liquid into which it was dipped is an acid, and a fortiori that it is not a non-acid. I do not need to know that the channel is working, and I do not need information about channel conditions. I only need to know that the liquid is an acid. So the drawn logical inference is sound, and so known. Thus in neither of these latter two scenarios is closure violated.

What we see is that even though closure holds here, according to the theory, it does not allow us to infer the worrisome claim that Adams et al. entertain that learning that something is an acid allows me to learn that channel conditions are in place. Closure does not tell me that no non-acid anywhere in the universe might turn litmus paper pink. That is because the fact that this liquid is an acid does not entail that there is not such a liquid *even though* I know it is an acid on the basis of the litmus paper turning pink. But what I *do* know, based on the fact that it is indeed an acid is that it is not a non-acid. My knowledge that being an acid entails not being a non-acid is not dependent the functioning of my local information channel. Rather the functioning of that channel is what allowed me to

[19] Recall we are here assuming that there is such a mystery substance somewhere in the universe.

get the information that I use as the antecedent for an entailment I already know.[20] The example shows nothing about closure but rather illustrates that channel conditions are situation specific.

Adams has many other examples, all roughly of this sort: an intuitively plausible situation where apparently one would attribute knowledge of p and of $p \rightarrow q$ to some agent, and yet deny knowledge of q to that same agent. What all of these examples have in common is that they function correctly only on the intuitive level, and then only to those who share the intuition that closure fails. When analyzed using the theory itself, and taking note of the fact that information as defined in the theory does flow across logical connectives (unless such information is already possessed), the examples all fail. In cases where the inference fails to generate knowledge under closure, it turns out that we should *also* deny the original knowledge attribution.

Let us ground this even more securely by analyzing an example using the mathematical structure of the theory: suppose Magda is an optimal epistemic agent and is in the lab observing the properties of various substances. She dips a blue litmus strip in a beaker of liquid and notes that it turns pink. Suppose that, on the basis of that observation, Magda comes to know that the liquid is an acid. What can we conclude? On ITAK we can say this: (1) Magda's belief that the liquid is an acid is caused (or causally sustained) by the information that it is an acid; (2) Therefore, based on her background knowledge alone, the probability that the liquid is an acid is less than 1, while based on her background knowledge and the signal constituted by her registering the color of the strip, the probability that the liquid is an acid is 1; (3) If it is an acid it is not a non-acid; (4) Magda can easily draw the conclusion that it is a non-acid.

Ok, so suppose she does draw the conclusion and believes on its basis that the liquid is a not a non-acid. By inverting the reasoning above we see that she now knows that the liquid is not a non-acid: the probability in advance of the signal that the liquid is not a non-acid is less than 1, and with the signal it is one, so the signal bears the information that the liquid is not a non-acid. And she believes it on the basis of that information. Thus she knows it. Note that this argument does not rely directly on the closure principle. Rather it relies on the fact that the situation for Magda is one where the probability structure clearly

[20] I am not asserting that the communication channel *would* function were there such a strange liquid. If it did not, we simply would not know that things are acids when they turn litmus paper pink. But the example *assumes* that it is functioning, and what follows from that …follows.

shows that she meets the conditions for knowledge. That is it appeals only to the correct mathematical fact that *information* is closed (and information flow as well for cases where one does not already have the information that would be generated by drawing known to be correct logical inferences).

Adams et al. do appear to be aware that a basic application of the rules of ITAK would indicate that the account really is closed under known entailment (2012). They there introduce the novel notion of "extracting information" and suggest that while (1) information is closed under the logical operations, and (2) if $p \rightarrow q$ and signal S carries the information that p, then S carries the information that q, still (3) the agent, A, may not be able to extract the information that q. Thus A will not know. (Adams et al. 2012: 192) This is an extra condition that Adams et al. are adding to the conception of knowledge, beyond the tenets of ITAK. What justifies its addition?

Their argument here is, again, an appeal to a variety of intuitions about cases. In this instance their focus is on contrastive knowledge statements: I know these are cookies rather than papier mâché; I know this is an acid rather than a non-non-acid; I know this is a zebra rather than a painted mule. They say "S may know that something is a zebra and not a mule if and only if S is in a position to know that it is a zebra *rather than* a mule" (Adams et al. 2012: 193). Their main contention is that in some situations in order to know that something is a zebra and not a painted mule I need both a zebra detector and a non-painted-mule detector. That seems obvious. The further conclusion, that some zebra detectors are not non-painted-mule detectors, however, is precisely the point at issue in discussions of closure. Because the question always, is, "does the fact that I know that this is a zebra in conjunction with the fact that I know that being a zebra implies being a non-painted-mule position me to know that this is a non-painted-mule?" Whatever the answer here, it simply does not follow from ITAK that the answer is "no." ITAK alone tells us that a zebra detector that suffices to generate knowledge of the presence of zebras will not identify painted mules as zebras. So we have the detector. What Adams et al. require is some further argument that when I use a zebra detector and come to know that I am seeing a zebra, somehow I cannot become informed that it is not a painted mule.

Rather than providing such an argument, Adams et al. (2012) have shifted to a consideration of information *extraction* and its difficulty to explain how we can have zebra detectors that are not non-painted-mule detectors. Yet they have not properly established that there are such detectors. Without offering any arguments about the nature of information extraction, or why it has conditions different from information transmission, they suggest that one might well be able to extract from some signal the information that this is

a zebra but not that it is not a painted mule.(Adams et al. 2012: 193) This is an interesting suggestion. But to bring it to bear on the issue of closure requires both that we have an account of what it is for information to be extracted from a signal so that some of it can cause a belief and some other part *contained in it* cannot, and also that we can find out in various cases whether or not it has done so. Adams et al. offer neither of these, perhaps because they believe that it follows as a mere conceptual matter that information extraction obeys distinct logical rules from information transmission. However they speak as though they have evidence that knowledge fails to obey the closure principle. But they have not displayed evidence of that sort. Instead their evidence is all to the effect that our intuitions tell us that closure does fail in certain cases. And it is about precisely these intuitions that people disagree. Intuitions aside, however, they have not shown that ITAK violates closure. It doesn't.

4.4 Summing up

This section has focused on one question. "Does ITAK violate closure?" I have examined the arguments that it does, and concluded that they fail, and I have argued directly that it does not.

Dretske and Adams and his collaborators may be right that the correct characterization of empirical knowledge is one on which closure fails. And it may well thus be that obeying closure constitutes a bad making feature of any proposed characterization of knowledge. What is clear, however, is that neither of these issues bears on the question of whether ITAK itself violates closure. It does not.

5 ITAK does rule out the KK thesis and so does not refute skepticism, but allows a decisive rejoinder

Something that many epistemologists, and Dretske in particular, have noticed is that it is not the epistemologist's job to refute the skeptic. And indeed, that is impossible in any case. Skepticism is irrefutable, for skeptical scenarios are entirely consistent with all the perceptual features of our experience. The trick is to figure out how to keep the possibility of skepticism from interfering with our epistemic practice, to quarantine it somehow or other. Dretske, Nozick and others have thought to quarantine skepticism by disallowing unrestricted closure. This restriction can work, but it comes with significant costs. Dretske, in advocating for the failure of closure, has acknowledged these costs, but has argued that

they must be paid in order to avoid the encroach of skepticism. Again, in outline the idea is that if skepticism cannot be refuted and if unrestricted closure is allowed, then modus tollens refutes all knowledge claims, even those as banal as "I have hands." But that is too fast. At most what modus tollens gives is that we cannot refute the skeptical *possibility* that we do not have hands. And yet, so what? That is just to restate the claim that we cannot refute skepticism. From the other direction one might think that knowing that we have hands, coupled with closure, allows us to refute skepticism. But that also is too fast. At most we get to know that skepticism is false, given that we can know that each individual skeptical scenario is false. Does that refute skepticism? I do not think that it does. Consider: my knowing that I have hands does not refute your claim that I do not know it. To refute your claim I need to show (perhaps to a third party) that I do in fact know it. And on this version of knowledge, knowing that I have hands simply does not put me in a position to show that I know it.

What would it be to refute the skeptic? Some things are clear in this respect. Showing that skepticism is a logical impossibility would suffice. But what about merely showing it false? On some ways of understanding refutation that would do it. Showing every possible skeptical scenario not to be our situation would also do. For example a detective considering many hypotheses about some crime will refute them one by one as she discovers evidence that is incompatible with each of them. And that would work if there were some limited number of skeptical possibilities. What about *knowing it to be false*? This is an even weaker position than showing false, at least for externalists about knowledge. That is because externalists need not always be in a position to show what they know. I will claim simply that for a true externalist position, knowing that p does not suffice to refute ¬ p. This is not because I believe closure fails, and that therefore one cannot be sure that knowing that p allows one to know that ¬ p is false. Rather refuting ¬ p seems to require that all reasonable doubt about ¬ p has been eliminated. But there seem to be many things that I know where my knowing them does not eliminate the possibility that I may be wrong. Many times my epistemic position when I am knowing successfully is cognitively indistinguishable from a state where I am not knowing successfully. Given that, there is no easy path from my knowing that p to eliminating the possibility that, after all, ¬ p is true. It *is not* true, but I am not always in a position to show it.

5.1 The KK thesis and channel conditions

In trying to explain why closure does fail after all, for his theory understood in the right way, Dretske says this:

> Our senses supply us with information about the world around us—that there are cookies in the jar or that our keys are on the table, for instance—thus enabling us to see whether there are cookies in the jar and where our keys are, but they do not supply us with all (some, but not all) the information implied by what they inform us about—that, for instance, there is a real physical world distinct from our impressions of it. That is why you cannot see that they (what you take to be cookies in the jar or your keys on the table) are not figments of your own imagination despite being able to see something (that they are cookies, that they are your keys) that implies they are not such figments. (Dretske, 2006: 411)

This is exactly right and exactly wrong at the same time. We do not, as Dretske says, get information about everything implied by *their informing us*. It is true that merely knowing something because a channel has brought us an information bearing signal about it does not suffice to inform us that the channel is working, for example. The fact that the source has the features it does is what the signal carries to us by way of the channel, but that fact is all that is necessary to give us whatever information entailed by *that fact*. So it does not matter for such an entailment that knowing on the basis of a channel does not tell us is that the channel itself is working. That the channel *is* working is implied by its informing us—only working channels inform—but *that we are informed* does not inform us that we are informed. This is a key distinction missed in every discussion of closure and the status of channel conditions. The reason it is so easy to overlook is that it operates *not* at the level of closure but instead at the level of the KK thesis.

Dretske himself has circled around this issue for many years, starting at least as far back as "Conclusive Reasons" where he identifies a "tendency to conflate knowing that P with knowing that one knows that P" (1971: 16). This is not news, but he goes on to correctly point out that the "confusion is partially fostered …by a failure to distinguish between what is implied in knowing that P, and what is implied (in some sense) by someone's *saying* he knows that P" (16). On this key confusion, however, Dretske (2004) fares no better than others. There Dretske is considering the case of Clyde, who sees some oranges and comes to know thereby that the are oranges. Clyde confidently asserts "I see some oranges." When asked whether he knows that they are not wax, Clyde admits that he does not. Applying closure he then finds that he does not know that they are oranges. At least this is Dretske's

argument for why we need to abandon closure: "Skeptical questions about whether you can tell oranges from wax imitations destroys your knowledge that they are oranges if, as closure tells us, you have to know (or be in a position to know) that they are not wax imitations in order to know they are oranges" (2004: 182). But this *at least* commits an error of equivocation. First Clyde should not say that he does not know that the oranges are not wax. Instead he should say that he does not know whether he knows that. In fact if he knows that they are oranges he is in a good position to know they are not wax, since being a real orange entails not being wax. By implicitly connecting Clyde's state of knowledge of the oranges with his state of knowledge about what he knows, Dretske comes dangerously close to making the very mistake he identifies (conflating knowing with knowing that one knows) and clearly conflates closure with the KK thesis itself, for only that conflation would support a move from not knowing that [rather whether] he knows they are not wax to not knowing that they are oranges. Can Clyde see that these are oranges, and thereby know that they are? Yes, by stipulation. Can he see that they are not wax? I have been arguing that he can, but Dretske wants to block this by denying closure. Is Clyde in any position to assert that he knows they are not wax? No. Why not? Because he is in no position to say that the signal indicating that they are not wax is information-bearing. But by the same token he is in no position to assert that he knows they are oranges. For by stipulation he cannot assert that the signal indicating *that he knows* is information bearing either. Being unsure whether he knows the oranges are wax requires being unsure about the goodness of the channel that is carrying signals about the oranges. Does this destroy Clyde's knowledge that they are oranges? No, at least not on ITAK. Nor does it affect his knowledge that they are not wax. Dretske was very close to seeing this, to seeing that it is the KK thesis that does all the work in these cases, but instead doubled down on anti-closure. His conclusion is that we can see that the things in front of us are the things they are, but not that they are things (rather than figments or other skeptical possibilities).

The thing is, you *can* see that the things in front of you are not figments of your imagination, if they are not. That is what seeing is. On the other hand you cannot use your seeing to find out that seeing really is a source of information about a real world. Only an independent channel carrying such information *about* the visual channel could do that. So here is what is required for knowledge that we have knowledge on ITAK, for example that we know *that we know that p*. If we know that we know that p then first p must be true, and also it must be true that we know that p. On ITAK to know that p requires that we believe p based on the information that p, and for that there must have been a signal bearing the information that p that bore that information to us. If we know that we know that p, then,

we know that there is such a signal. That knowledge then does commit us to knowing that there is (or was at the time) a functioning communication channel that carried a signal with that information. Notice that we do not need to know any of this when we merely know that p; we only need to believe that p based on the information that p. ITAK makes no version of the KK thesis remotely plausible.

Here now is one suggestion about what it would take to rule out ¬ p: only when I am in a position to know that I know that p am I in a position to rule out ¬ p. If we adopt such a view of ruling out, we will see that ITAK does not rule out skepticism, despite the fact that closure holds.

Recall that channel conditions are those conditions that allow for a communication channel:

> The channel of communication = that set of existing conditions (on which the signal depends) that either (1) generate no (relevant) information, or (2) generate only redundant information (from the point of view of the receiver).

Generally, as Dretske says, none of these conditions can be known merely because we know something or other based on the signal the channel passes. But as we showed above that is because the signal does not, in fact, bear information about the channel conditions. We also showed in that context that there is no need to construct a failure of closure to account for this feature of communication channels. In most cases we can, of course, learn that the channel conditions hold. But we do so by calibrating the channel in Sense 2, that is by taking features of the world that we know about independently of the channel, and checking that the channel would properly pass information about those features.

I now make the (I hope) uncontroversial assumption that we are all empiricists of some stripe or other. At least I will rely on this much empiricism: all empirical knowledge arises as the result of sensory experience that gives us knowledge by informing us about the world. Now let us ask whether, given that we know something about the world—for example that we have hands—we can rule out skeptical hypotheses, where ruling out requires, as I argued above, knowing that we know such hypotheses are false.

To know anything, on ITAK, requires believing it on the basis of an information-bearing signal. To know that I know any empirical fact at all (like that I have hands) then requires a signal to flow that bears the information that I know that I have hands. There must then be a signal that bears the information that I believe that I have hands *on the basis*

of an information-bearing signal.[21] To know of some signal that it bears information is to know that it travels in a properly functioning communication channel. So finally we get to the crux of the matter: can we know that empirical communication channels function properly? The surprising answer is that while we can know of individual empirical channels that they are properly functioning, *given that* we are not in a skeptical scenario, we *cannot* know that empirical channels as such can carry information. For what signal could bear the information that empirical channels as such are suited to carrying information bearing signals? In principle such a thing is impossible; it amounts to a signal bearing information to the effect that it is bearing that information. So if skepticism is false ITAK allows a great deal of empirical knowledge, as one would hope. But even if skepticism is false I cannot know that I know it, and therefore I cannot refute the skeptic.

One might worry at this point that I could use my ears and other senses to get the knowledge that my eyes are functioning information channels, and then I would know that I know that skepticism is false by first coming to know that I have hands by using my eyes, and then (using my other senses) come to know that my belief that I know that I have hands was caused by an information bearing signal (transmitted through my eyes).[22] However unlike the cases considered above in the discussion of closure, this really would amount to an illegitimate appeal to various channel conditions being in place. To make this case more clearly, I will now consider the important and tricky notion, calibration, and then apply the machinery of ITAK to evaluate whether we can know that we know that skepticism is false (under the assumption that we know that it is false). We will find that we cannot.

Recall our discussion of calibration above. To know that I know that p requires calibrating in Sense 2 the channel that bears the information that p. To do that requires that I get the information that the channel is calibrated via a distinct communication channel that is itself capable of bearing that information. And *that* puts significant demands on the new channel. Suppose the original channel involves at some stage the transmission of electricity through a copper wire network. If the wires are hidden and what we need to check is whether they are properly connected, we can use a communication channel

---

[21] Remember, that is just what knowledge claims amount to on ITAK.

[22] Clearly my visual sense comprises many individual (and overlapping) information channels. The fact that I can use their gross functioning, say, to check their fine functioning (as I do using an eye chart) introduces no special complications beyond this fact. In the same way my WiFi has two channels one of which can check the other (when it is functioning) even though in most contexts I will treat it as a single communication channel.

involving the right configuration of copper wires to get that information by, say, connecting it at the end points of the other wire. But what we cannot do is use a device whose operation is predicated on the fact that copper wires obey Ohm's law to determine that relying on Ohm's law works. An Ohm's law based device, whose use presupposes rather than tests Ohm's law, is incapable of transmitting the information that Ohm's law works.

Now apply these considerations to our senses understood as channels that carry various sorts of signals: visual, auditory, tactile, etc. Each sensory modality can be used in the right circumstance to provide a check on the others. In viewing a hologram, my eyes are not well-calibrated to give me information about 3-dimensionality, and by passing my hand through the hologram I can find this out tactually. Indeed I can even use my eyes to observe features of my eyes that may be preventing them from working properly in some contexts, for example observing how much flexion there is in the lens given certain external prompts may let me know that I have become farsighted. Neither of these, nor a myriad of other calibrating uses of my senses requires that they transmit the kind of information about their own channel conditions that are relevant to the transmission of information about others' channel conditions, and so they do not run afoul of the constitutive structure of ITAK. We really can come to know that we know that our senses are well-calibrated. But the sense of calibration here is sharply restricted. We can know that they are well-calibrated only with respect to *other* channels of their type. What type is that? The empirical type. I can find out a lot about various empirical channels of communication. But what I cannot find out empirically is whether these empirical channels are themselves connected to the world as opposed to vat signals, or the illusions of some clever, malicious and powerful spirit. For in such a case every signal would be indistinguishable from the case where the world is largely the way our experience tells us it is.

This result shows that ITAK provides precisely the option that Dretske and others envisioned allowing us to have empirical knowledge but not allowing us to rule out skepticism. Moreover, properly understood ITAK does this without tying our inferential hands by restricting closure. One may well argue at this point that ITAK is itself unsound for other reasons. I do not think that it is, but it has not been my purpose to make that case. Rather my aim here has been to reassure those who agree with Dretske that skepticism cannot be ruled out, and who find the basic structure of ITAK compelling in itself, that they need

not abandon closure. That I think I have done.[23]

James Mattingly
Georgetown University
Philosophy Department,
3700 O Street NW
Washington, DC 20057
jmm67@georgetown.edu

References

Adams, Fred, John A., Barker and Figurelli, Julia. 2012. "Towards closure on closure". *Synthese 188*: 179–96.

Baumann, P. 2006. "Information, closure, and knowledge: on Jäger's objection to Dretske". *Erkenntnis 64*: 403–8.

Dretske, Fred. 1970. "Epistemic operators". *Journal of Philosophy 67*: 1007–23

Dretske, Fred. 1971. "Conclusive reasons". *Australasian Journal of Philosophy. 49*(1): 1–29

Dretske, Fred. 1981. *Knowledge and the Flow of Information*. Cambridge: The MIT Press.

Dretske, Fred. 2004. "Externalism and modest contextualism". *Erkenntnis 61*: 173–86.

Dretske, Fred. 2005. "The case against closure". In *Contemporary Debates in Epistemology*. Matthias Steup and Ernest Sosa (ed.). Malden, MA: Blackwell: 13–26.

Dretske, Fred. 2006. "Information and closure". *Erkenntnis 64*: 409–13.

Greco, John. 2010. *Achieving Knowledge*. Cambridge: Cambridge University Press.

Jäger, Christoph. 2004. "Skepticism, information, and closure: Dretske's theory of knowledge". *Erkenntnis 61*: 187–201.

Kipper, Jens. 2016. "Safety, closure, and the flow of information". *Erkenntnis 81*: 1109–26.

Mattingly, James. 2021. *Information and Experimental Knowledge*. Chicago: Chicago University Press.

Shackel, Nicholas. 2006. "Shutting Dretske's door". *Erkenntnis 64*: 393–401.

[23] Thanks are due to many people, but here I want to especially single out Ruth Marcela Espinosa, John Greco and Bryce Huebner for important feedback and support. I would like to acknowledge an anonymous referee as well as Elia Zardini for many useful comments.

Wright, Crispin. 2003. "Some reflections on the acquisition of warrant by inference". In *New Essays on Semantic Externalism and Self-Knowledge*. Susanna Nuccetelli (ed.). Cambridge: The MIT Press.