Knowledge-Closure and Skepticism (J)

> Uncovering difficulties in the details of particular formulations of [the closure principle] will not weaken the principle's intuitive appeal; such quibbling will seem at best like a wasp attacking a steamroller, at worst like an effort in bad faith to avoid being pulled along by the skeptic's argument. Robert Nozick, *Philosophical Explanations*, p. 206

Leopold enters his dining room and seems to see a chair in front of him. Does Leopold know that there is a chair in front of him? The following informal skeptical argument attempts to show that he does not:

> Leopold does not know that he is not a brain in a vat inhabiting a chairless world.
> If Leopold does not know *that*, then he does not know that there is a chair in front of him.
> So, Leopold does not know that there is a chair in front of him.

This argument is valid. Its first premise might be defended on the grounds that the particular skeptical hypothesis involved—the hypothesis that he is a brain in a vat inhabiting a chairless world—is well-chosen so that Leopold does not know it to be false. Alternatively, it might be defended on the general grounds that agents like Leopold cannot possibly know any skeptical hypotheses to be false. In this paper we will not have much to say about this first premise. We will mostly be concerned with the argument's second premise and with some issues pertaining to the argument's conclusion.

The second premise of the argument might be defended by invoking some principle of *epistemic closure*. Although there are other ways to defend the second premise, this type of defense seems to be the best known and most prominent one. We will discuss it at length. We intend to examine skeptical arguments that are based on epistemic closure principles.[1]

The conclusion of the argument is that one agent (Leopold) does not know one thing (that there is a chair in front of him). This conclusion is not immediately troubling. Many agents do not know many things. The argument is supposed to be worrisome because it seems obvious that a similar and equally plausible argument could be given for the conclusion that any typical agent does not know any typical proposition ordinarily thought to be known by the agent on perceptual grounds. In other words, the argument concerning Leopold and the chair is supposed to be worrisome because it is seen as an argument from arbitrary instance to the quite general conclusion that all human agents lack all perceptual knowledge. We will pay some attention to the details of how the generalization from Leopold and his chair to human agents and perceptual knowledge in general might go.

We begin in Section I with some preliminary remarks about epistemic closure principles, knowledge-closure principles in particular, and the role of knowledge-closure principles in skeptical argumentation. After discussing some implausible knowledge-closure principles and identifying two key problems for such principles in Section II, we move on in Sections III and IV to discuss more plausible knowledge-closure principles and their possible use in skeptical argumentation. We conclude that knowledge-closure based skeptical arguments are not promising skeptical arguments.

I. Preliminaries

An epistemic closure principle is a principle asserting that if an agent stands in some epistemic relation, say the relation of knowing, to a proposition, then the agent also stands in that epistemic relation to further propositions connected in some specified way to the first proposition. The epistemic relation is then said to be "closed" under the specified connection. For example, the following principle expresses the closure of knowledge under entailment (i.e., necessary consequence, strict implication):

Closure 1: For all *S*, all *p*, all *q*: if *S* knows *p*, and *p* entails *q*, then *S* knows *q*.

This says that, for all agents *S*, and propositions *p* and *q*, if *S* knows *p*, and if *p* entails *q*, then *S* knows *q*. Note that epistemic principles, like philosophical principles in general, are traditionally intended as necessary truths. So the principle should be understood to fail unless it holds with necessity. From now on, we will sometimes follow general custom and suppress initial quantifiers as well as references to the agent; we will also use obvious symbols and insert clarifying parentheses. The short version of Closure 1 will then look like this:

Closure 1: (K*p* and (*p* entails *q*)) $\rightarrow$ K*q*.

Some (but not all) skeptical arguments employ a knowledge-closure principle. Here is an example of a skeptical argument that employs Closure 1. Actually, it is just an argument-sketch representative of how closure-based arguments are frequently portrayed in the literature on closure and/or skepticism. Consider the following skeptical hypothesis: h = you are the victim of an evil demon who is leading you into falsely believing all of the things you take yourself to know via perception; e.g. the demon is misleading you into believing that there are chairs in the room when in fact there are not. Let c = the proposition that there are chairs in the room; let ~h = the denial of the skeptical hypothesis just sketched, i.e., the proposition that you are *not* the victim of an evil demon who… The skeptic then argues as follows:

P1. (K*p* and (*p* entails *q*)) $\rightarrow$ K*q*
but
P2. ~K~h
hence
3. Either ~Kc or ~(c entails ~h)
but
P4. c entails ~h
hence
5. ~Kc

Premise P1 is just Closure 1. Premise P2 claims that you don't know that the skeptical hypothesis is false (perhaps you can't, perhaps you simply don't). This is the skeptic's key claim—one might, of course, challenge this claim, but that would be a topic for another paper. Step 3 follows from P1 and P2. Premise P4 is supposed to be obvious. The conclusion, 5, follows from 3 and P4. If sound, an argument along these lines might be taken to show that agents such as us don't know the humdrum propositions we take ourselves to know via perception, such as the proposition that there are chairs in the room.

The basic strategy behind this type of skeptical argument is this. First employ a general epistemic principle the non-skeptic is committed to, then use a skeptical hypothesis to turn the principle against the non-skeptic. It is crucial to the success of this strategy that the non-skeptic really be committed (at least implicitly) to the general epistemic principle in question. This is *not* the case for the sample argument given above (we mention it for purposes of illustration only). The simple closure principle it employs, Closure 1, is easily seen to be false: people do not know all the necessary consequences of what they know. The argument, though valid, is unsound. But perhaps there is a more promising closure-based skeptical argument around the corner. Such an argument would have to employ a closure principle that is not easily seen to be false, and it would have to lead (through at least arguably sound steps) to an interesting skeptical conclusion.

Here we will focus on closure principles involving one epistemic notion: *knowledge*. We will not discuss closure principles involving justification, or warrant, or rationality, or other epistemic notions. We don't mean to suggest that discussing such closure principles would not be worthwhile, or that such a discussion would not be of relevance to skeptical argumentation. We focus on closure principles involving knowledge because most discussions of closure-based skepticism turn on closure principles involving knowledge, and also because we like knowledge and want to think about knowledge and knowledge-skepticism. Also, we tend to think that knowledge is somewhat easier to handle than notions like justification or rationality. The latter are ambiguous in ways in which knowledge is not (if the notion of knowledge is ambiguous, it is less ambiguous than the other notions).

Our central question is this: Are there any plausibly true epistemic closure principles concerning knowledge that figure centrally in arguably sound skeptical arguments with interestingly strong conclusions? There are several angles from which one might approach this question. One might first try to formulate a plausible skeptical argument and then ask what closure principle, if any, the argument employs. Alternatively, one might look directly at closure principles, asking of each candidate principle whether it is vulnerable to counterexample; if not, one might go on, asking whether a skeptical argument using the principle looks promising. We will adopt this latter approach.

Most (but again, not all) contemporary discussions of skepticism involve knowledge-closure principles, so it seems that quite a few epistemologists believe that the answer to our question is: "Yes, there are indeed plausible knowledge-closure principles that figure centrally in arguably sound skeptical arguments with interestingly strong conclusions." It turns out, however, that authors whose main concern is with skepticism often invoke implausible closure principles when discussing closure-based skeptical arguments. Sometimes they even invoke a principle they themselves acknowledge to be false, adding a footnote to the effect that the principle would have to be modified somehow to make it immune to counterexample. On the other hand, authors who explicitly discuss closure principles tend to focus on providing counterexamples against some such principles or on protecting them from alleged counterexamples. Their discussions are usually motivated by a background concern with skepticism, but they almost never include even minimally detailed consideration of how, if at all, the most plausible closure principles might figure in a specific skeptical argument. We expect that others who are familiar with the relevant literature have occasionally wondered about this situation and have thought that someone should try to sort out what exactly is going on with knowledge-closure principles and skepticism.

It is, after all, possible to smell a rat here. One might reasonably suspect that there are not that many true closure principles at all and that the few plausible ones will be of less use to the skeptic than it initially appeared while the ones that would clearly be of use to the skeptic will not be plausible. This suspicion motivates our discussion. Though we are inclined to agree with Feldman (1995, 487) that "some version of the closure principle…is surely true", this does not force us to believe that the true version(s) of the principle will be useful in a skeptical argument. This is an issue that depends on the details. So, *pace* Nozick (see the passage with which we began this paper), we intend to quibble about the details of closure principles because the details may well matter to the formulation and evaluation of a specific skeptical argument.[2] Though a few authors have looked into some of the relevant issues,[3] the main issues that concern us here have not been addressed in sufficient detail and have certainly not been resolved. We don't promise to resolve them here, but we do intend to take some steps towards a resolution.

II. Implausible Principles

Closure 1, recall, claimed that knowledge is closed under entailment (necessary consequence); that is, it claimed that one knows the necessary consequences of what one knows:

Closure 1: (K*p* and (*p* entails *q*)) $\rightarrow$ K*q*.

This principle is false for at least two reasons. First, and most obviously, the principle implies that you believe all the necessary consequences of what you know, and this is clearly implausible. Vast numbers of propositions are entailed by what you know, far too many to believe them all. Second, assume you *do* believe *q*; and assume, in addition, that *q* is a necessary truth. Necessary truths are trivially entailed by any proposition. So, according to Closure 1, you are guaranteed to know *q* as long as you know a single truth. Surely, that's much too lenient: you might believe *q* for bad reasons; or you might believe it in the face of strong (albeit misleading) evidence to the contrary, evidence that you are unable to defeat. If you did this, you would not know *q* even though it is entailed by something you know. Closure 1 is invitingly simple but false: it cannot serve in a successful skeptical argument.

We have just seen that Closure 1 suffers from at least two problems. It seems natural to begin to respond to these problems by adding a requirement to the effect that the agent must have at least "noticed" the relevant entailment. This leads to the following well-known closure principle saying that knowledge is closed under *known* entailment:

Closure 2: (K*p* and K(*p* entails *q*)) $\rightarrow$ K*q*.

The skeptical argument employing Closure 2 goes roughly like this (where again c = some bit of alleged perceptual knowledge and ~h = the denial of the skeptical hypothesis):

P1. (K*p* and K(*p* entails *q*)) $\rightarrow$ K*q*
P2. ~K~h
3. ~Kc or ~K(c entails ~h)
P4. K(c entails ~h)
5. ~Kc

Although Closure 2 is the best known and most widely discussed epistemic closure principle, it is still false. It is the principle one usually finds in discussions of closure-based skepticism, sometimes with an accompanying footnote indicating that it would have to be modified somehow to make it immune to counterexample.[4]

Closure 2 faces two problems. Here is the first problem. One can fail to believe *q*, despite knowing *p* and knowing that *p* entails *q*. One might not have "put two and two together" and done the inference (knowledge of an entailment does not guarantee that one actually draws the inferences it underwrites), and one might not otherwise believe *q*. Call this the "belief problem".

Here is the second problem. We call it the "warrant problem". Let us bracket the belief problem and pretend agents who know *p* and know that *p* entails *q* also believe *q*. Closure 2 requires not merely that *q* be believed; it requires that *q* be *known*. Consistent with knowing *p* and knowing that *p* entails *q*, however, one might believe *q* for terrible reasons (and perhaps in the presence of relevant defeaters, etc.). While the belief problem raises the psychological issue whether we will find all the beliefs Closure 2 predicts we will find, the warrant problem raises the epistemic issue whether such beliefs, even where they are found, will have the required epistemic status to count as knowledge.

Before we turn to principles that directly address these problems we want to take a look at some relatives of Closure 2 which we will ignore in our subsequent discussion. Feldman (1995, 488) discusses a principle about the "transmission of justification" which, reformulated as a principle about the "transmission of knowledge", would look like this:

Closure 3: $K(p \text{ and } (p \text{ entails } q)) \rightarrow Kq$.

Evidently, Closure 3 does not solve the belief problem; and it does not solve the warrant problem either. Still, there might be something to be said for it. Note first that Closure 3 should not be any worse off than Closure 2, for it is plausibly taken to be entailed by Closure 2.[5] In addition, it looks like Closure 3 ought to be a bit safer than Closure 2, for Closure 3 does not obviously entail Closure 2. Now consider the rule

R1: $Kp \text{ and } K(p \text{ entails } q) \rightarrow K(p \text{ and } (p \text{ entails } q))$.

Rule R1 is a restriction of a more general rule, the so-called "conjunction rule", according to which K*x* and K*y* together entail K(*x* and *y*), for any propositions *x* and *y*—a rule that figures prominently in various versions of the lottery paradox and has been contested.[6] Since R1 is a restriction of this contested conjunction rule, R1 itself might be considered questionable. Now, though Closure 3 does not obviously entail Closure 2, the combination of R1 with Closure 3 does obviously entail Closure 2. In view of this, one might suspect Closure 2 of implicitly relying in some manner on the questionable R1. If this is a serious worry for Closure 2—we are not sure whether it is—Closure 3 will be preferable: although it does not begin to address the belief problem and does not do much by way of addressing the warrant problem, it at least avoids a potential difficulty.

Closure 3 allows us to illustrate how modifications to closure principles may significantly affect their utility for skeptical argumentation. (This is a bit artificial because Closures 2 and 3 both face difficulties that suffice for their not being useful to the skeptic, but the illustration should prove useful nonetheless.) A skeptical argument employing Closure 3 would *begin* like this:

P1. $K(p \text{ and } (p \text{ entails } q)) \rightarrow Kq$
P2. ~K~h
3. ~K(c and (c entails ~h))
P4. K(c entails ~h)

But now it is unclear how the skeptic is to get from this point to the desired conclusion "~Kc". An epistemic principle underwriting the step from 3 and P4 to "~Kc" would be:

P5. ~K(*p* and (*p* entails *q*)) → ~K*p* or ~K(*p* entails *q*)

which, together with 3, allows the skeptic to derive the disjunction "~Kc or ~K(c entails ~h)", which in turn gets him to "~Kc" by premise P4.

But principle P5 is logically equivalent to the rule R1. So the skeptic seems to weaken his argument, making it vulnerable to the charge of relying on a questionable rule. Moreover, if a skeptical argument based on Closure 3 relies on R1 anyway, then the switch from Closure 2 to Closure 3 is ultimately pointless: with R1 available, the sole ground for suspicion with Closure 2 that motivated switching to Closure 3 has been removed (since R1 and Closure 3 together entail Closure 2, any worry about the latter is a worry about the conjunction of the former). To make a strong argument, the closure-based skeptic needs the best closure principle he can get. Since Closure 3, *qua* closure principle, appears to be somewhat safer than Closure 2, the skeptic might consider it to be a preferable starting point. But in order to employ it in his argument, he needs to rely on a vulnerable rule combined with which Closure 3 offers no improvement over Closure 2 for skeptical purposes. The skeptic might as well have stuck to Closure 2: the net gain in plausibility of the argument based on Closure 3 seems to be zero.

In what follows, we will ignore variants of closure principles which, in the manner of Closure 3, combine two occurrences of the K-operator in the antecedent into one. Though they might be somewhat safer than their "distributed" siblings, this gain in safety will be offset by the need to complicate the skeptical argument when employing such a principle.

Let us briefly look at another relative of Closure 2. We have noted that Closure 2 is the most prominent knowledge-closure principle by far. We have also noted that it is typically introduced as an improvement over Closure 1. But one might well wonder what exactly motivates moving from Closure 1 to Closure 2, rather than

Closure 4: (K*p* and (*p* entails *q*) and B(*p* entails *q*)) → K*q*,

which would seem to be an alternative way of accommodating the idea that in order to know *q*, when *q* is entailed by something one knows, one must have "noticed" the relevant entailment. Why is Closure 2 so prominent while Closure 4 is never mentioned? Sure enough, Closure 4 suffers from the belief problem (and the warrant problem); but so does Closure 2. Is there anything Closure 2 can do that Closure 4 can't do?

Here are three considerations. First, say an agent knows *p* and believes correctly that *p* entails *q* but doesn't know this because the agent holds the entailment belief based on hearsay rather than understanding and/or reflection. Such an agent might be said to fail to know *q* on the grounds that she will be unable to defeat (misleading) evidence against (*p* entails *q*) and/or against *q*. An agent who knows *p* and *knows* that it entails *q*, on the other hand, should be better positioned (in principle) to defeat such misleading evidence; so Closure 2 seems better than Closure 4. Second, what exactly is it to believe that *p* entails *q*? Say some agent believes that *p* entails *q* but doesn't believe that, necessarily, if *p* is true then *q* is true. Is this possible? Well, why not? Does the agent know *q*? One might think she does not have sufficient grasp of the concept of entailment to really know *q* given what she has to go by. As a remedy one might propose this: though it's possible to believe that *p* entails *q* while having a defective grasp of the concept of entailment, it's not possible to *know* that *p* entails *q* while having a defective grasp of the concept of entailment; so Closure 2 once again seems better than Closure 4. Third, Closure 2 does not appear to entail Closure 4, whereas Closure 4 trivially entails Closure 2.[7] So Closure 2 appears to be the safer principle. In what follows, we join the general custom of neglecting Closure 4.

III. Plausible Principles (a): Strengthening the Antecedent

We now want to return to the main line of the paper. Closures 1 to 4 share a consequent, namely K*q,* that gets the principles into trouble with the belief problem. There is nothing in the antecedents of any of these four principles guaranteeing that *q* is believed, and without believing *q* one doesn't know *q*. Searching for a closure principle that is not clearly false, one can pursue two general strategies. One can (a) build belief in *q* into the antecedent of a candidate closure principle, or (b) weaken the consequent of a principle so that it does not require belief in *q*. We will now explore these moves in some detail. In this section we consider option (a). In section IV we consider option (b).

Let us first consider a simple principle that aims to avoid the belief problem by building belief in *q* into the antecedent of Closure 2:

Closure 5: (K*p* and K(*p* entails *q*) and B*q*) $\rightarrow$ K*q*

Though Closure 5 indeed sidesteps the belief problem, it does nothing to address the warrant problem: it says nothing at all about one's reasons for believing *q*. What if, for example, one knows *p*, knows that *p* entails *q*, and believes *q*, but believes *q* not because of *p* and the entailment but rather for some problematic reason? For example, what if I know that there are at least 9 people in the room (I started counting and stopped at nine), and I know that this entails that there are at least 7 people in the room, and I believe that there are at least 7 people in the room but not because of the inference or the counting but rather because I *always* believe that there are at least 7 people in the room when I'm in the room. Though the example is somewhat frivolous, it should be clear that in this sort of situation, where one does not properly base believing *q* on proper grounds for believing *q*, it looks like one does not know *q*.

The obvious response is to build an appropriate basing relation clause into the antecedent of Closure 5. The result will be a principle similar to:

Closure 6: (K*p* and K(*p* entails *q*) and B*q* based on deduction from *p* and (*p* entails *q*)) $\rightarrow$ K*q*

We will not offer a detailed exploration of various difficulties concerning the nature of the basing relation. But a few remarks seem to be in order.

It should be noted that the inserted clause about the basing relation, "B*q* based on deduction from *p* and (*p* entails *q*)", is a bit ambiguous. It can be understood in at least two different ways, reflecting two views about deduction and about what sorts of beliefs Closure 6 expects the epistemic agent to hold. More precisely, the clause can be understood to imply either (i) or (ii):

(i) *S* believes *q*, and *S* believes *p*, and S believes that (*p* entails *q*), and *S* believes *q* based on deducing it from her belief *p* together with her belief that (*p* entails *q*).

(ii) *S* believes *q*, and *S* believes that (*p* and (*p* entails *q*)), and *S* believes *q* based on deducing it from her belief that (*p* and (*p* entails *q*)).

Reading (ii) demands, as it were, more by way of beliefs from the agent than what is already involved in K*p* and K(*p* entails *q*). It expects the agent to have put together into one conjunctive belief the two beliefs involved in these two bits of knowledge. Accordingly, it assumes a picture of deduction where the agent deduces her belief *q* from her conjunctive premise-belief that (*p* and (*p* entails *q*)). Reading (i) demands less from the agent. It does not expect that she has put together the beliefs involved in K*p* and K(*p* entails *q*) into a single conjunctive premise-belief. Accordingly, it assumes a picture of deduction where the agent deduces her belief *q* from two premise-beliefs, without expecting her to have put the two together into a single premise-belief. On the face of it, reading (i) seems preferable: since it expects less from the agent, the resulting version of Closure 6 will apply to a wider range of cases than a version incorporating reading (ii). We therefore select the first reading.[8]

All the principles preceding Closure 6 are easily understood as making synchronic claims, saying that an agent possesses as piece of knowledge *q* at a time *t* provided the agent satisfies a certain condition at that same time *t*. Although Closure 6 as a whole should also be understood in this way, it must be noted that the clause about the basing relation complicates matters because it carries with it an implied reference to a time earlier than *t*. An agent who is at a time *t* in the state of believing *q based on deduction* from premises *p* and (*p* entails *q*) will typically have entered that state by way of an inference that began at a time prior to *t* at which the agent did not yet believe *q*. Moreover, the basing clause should be understood to imply that the agent already knew the premises at that prior time, otherwise the clause would probably not remove the problems it was designed to remove.

We note, finally, that Gilbert Harman has argued repeatedly (e.g. 1973, chap. 10) that it is at best misleading to talk about deduction as if it were a psychological inference process. We do not wish to engage this issue here. We assume that Closure 6 can be interpreted (or, if need be, reformulated) so as to avoid any mistaken views about the nature of deduction.

Some have thought that Closure 6 is beyond criticism: Doesn't it merely say that one knows what one (properly) deduces from what one knows? And isn't this so obvious as to be beyond criticism? Well, that's two questions. The answer to the first is "Yes—although, to be more precise, Closure 6 actually says that one knows what one has (properly) deduced from what one knew and still knows." The answer to the second question is "No, this is not so obvious as to be beyond criticism."

Closure 6 may run into trouble with cases of epistemically *overdetermined* beliefs. I know that this is a chair, and I know that if this is a chair it cannot fly away, and I believe on the basis of these two bits of knowledge that the chair will not fly away. Additionally, however, I believe that the chair will not fly away because I believe that its wings are broken (and also because I believe that my crystal ball revealed to me that it will not fly away).[9] Admittedly, this is again a frivolous example. Perhaps the following is more realistic. Pretend I'm a thorough logician and, whenever I can manage, I like to provide two independent proofs of every theorem I prove. On one occasion I knowingly deduce a theorem, T, from known starting assumptions but also independently manage to fallaciously "deduce" it from misunderstood starting assumptions. Do I know T? And what if I used additional bad proofs?

We are not quite sure what to make of cases like this—much seems to hang on exactly how one understands the difficult notion of believing one thing "on the basis of" some other things. Much seems to hang, that is, on further detailed issues about "the basing relation" left unsettled by our brief discussion earlier. Perhaps whether one knows depends in part on relevant counterfactuals expressing what one would believe if the other reason(s) were subtracted out and/or on the firmness of belief which comes from each individual source of belief. It may well be worth exploring these issues about overdetermination, closure, and the basing relation in some detail.[10] But when searching for a plausible closure principle, it is perhaps best to avoid such complexities and consider a slightly modified candidate:

Closure 7: (K*p* and K(*p* entails *q*) and B*q* *solely* based on deduction from *p* and (*p* entails *q*)) $\rightarrow$ K*q*.

We think there are challenges even to this extremely plausible looking principle. Here we will merely gesture at three of them: There might be agents with weird global defeaters; in particular, agents who satisfy the antecedent of the principle even though they believe (and for subjectively quite compelling reasons) that deduction is not a good way of adding beliefs.[11] There might incoherent agents; in particular, agents who satisfy the antecedent, hence believe *q*, but also incoherently believe ~*q*: Can one really know *q* while believing ~*q*? And there is also a quite different challenge deriving from cases caught up in epistemic paradox.[12] Though we take these challenges quite seriously, we acknowledge that they involve relatively arcane issues. We will not pursue them at this point.

There are principles intuitively "in between" Closure 6 and Closure 7 that might also be of interest. An example of such a principle is:

Closure 6.5: (K*p* and K(*p* entails *q*) and (B*q* based on deduction from *p* and (*p* entails *q*)) and ~(B*q* based on some "defective reason")) $\rightarrow$ K*q*

If one interprets the "no defective reason"-clause broadly enough, this principle avoids the "bad" overdetermination cases facing Closure 6 and allows for more flexibility in the antecedent than is allowed by the antecedent of Closure 7. On the downside, because of the "no defective reason"-clause, Closure 6.5 is rather more vague than Closure 6 and Closure 7. Still, it might be considered preferable to both. Note, however, that the challenges to Closure 7, briefly hinted at in the previous paragraph, will be equally relevant to the evaluation of Closure 6.5—and, of course, to Closure 6 as well.

We have now identified three closure principles, Closures 6, 6.5, and 7, that appear to be sufficiently plausible to make it worth examining the question whether they are of use in a skeptical argument. Though Closure 6 is rather more vulnerable than Closures 7 and 6.5 (which are themselves not beyond challenge), it is not obviously false. Since it is also the least cumbersome of the three, we will begin by asking how it performs in a closure-based skeptical argument.

Let us first look at an argument-sketch starting from Closure 6 that proceeds along the lines of the argument sketches given earlier—we abbreviate the cumbersome locution "based on deduction from" to "bod from":

Argument Sketch:

P1. K*p* and K(*p* entails *q*) and B*q* bod from *p* and (*p* entails *q*) $\rightarrow$ K*q*

P2. ~K~h
3. ~Kc or ~K(c entails ~h) or ~(B~h bod from c and (c entails ~h))
P4. K(c entails ~h)
P5. B~h bod from c and (c entails ~h)
~Kc

This way of representing the skeptical argument, though useful for indicating the argument's overall structure, has some shortcomings. Most importantly, rendering the conclusion as "~Kc"—where "c" is said to stand for some bit of alleged perceptual knowledge and the reference to agents is entirely suppressed—tends to downplay the conclusion the skeptic ought to be aiming for. When the "specter of skepticism" is raised, the skeptical threat is described as being of a very general nature. Skeptical arguments are supposed to threaten us with the strong general conclusion that we (i.e., all human agents) lack all knowledge of an interesting sort (all empirical knowledge, or all perceptual knowledge); they are supposed to show that we don't know the empirical or perceptual propositions we think we know. Thus, the conclusion the skeptic really ought to be aiming for must say that no human agents know any proposition of some interesting class C:

7. $\forall S, \forall p \in \mathrm{C}$: ~$S\mathrm{K}p$,

For present purposes we should think of class C, the class of propositions the skeptic is targeting, as the class of "perceptual propositions"—where perceptual propositions are, very roughly, propositions about our (alleged) environment (including our own bodies) which we typically believe or disbelieve (if we believe or disbelieve them at all) based, at least in part, on how our (alleged) environment looks, sounds, tastes, smells or feels to us. Whether this class can actually be specified in a more precise and satisfying manner is a difficult issue. We will sidestep it for now, noting that the task of specifying C will fall to the skeptic.[13]

The skeptic under consideration here operates with a skeptical hypothesis, say some version of the evil-demon hypothesis, which is extracted from a larger skeptical scenario—a scenario that "explains", through the machinations of some evil demon, how it might have come about that we hold our perceptual beliefs, based on how things perceptually appear to us, even though these beliefs are all false. Skeptical hypotheses extracted from such a scenario are designed to attack our *perceptual* beliefs. They are not designed to attack beliefs that are concerned exclusively with our own conscious states or with simple logical matters. This accounts for the restriction to propositions of class C. Moreover, such skeptical hypotheses are designed to attack our perceptual *beliefs*. They are not designed to attack propositions we do not believe, for they offer an explanation of why we hold the perceptual beliefs we hold compatible with their being false. So the skeptic will try to reach his ultimate conclusion, 7, via a slightly weaker lemma,

6. $\forall S, \forall p \in \mathrm{C}$: $S\mathrm{B}p \rightarrow$ ~$S\mathrm{K}p$,

which says that no human agents know any propositions of class C they believe. The more general conclusion, 7, follows easily from 6 by the principle that no one knows anything they don't believe.

Having identified the conclusion the skeptic ought to be aiming for, we should be able to turn the above argument-sketch into a more carefully spelled out general argument. It is usually taken for granted that this involves interpreting the argument sketch as attempting to establish the general skeptical conclusion, 7, by way of arguing from “arbitrary instances” (arbitrary human agents and arbitrary perceptual beliefs held by such agents). Before we look and see how this is supposed to go, we must address an additional wrinkle that comes up with respect to the fourth step of the argument sketch.

Premise P4 of the argument-sketch says “K(c entails ~h)”. Since knowledge entails truth, this premise requires that it must be true that (c entails ~h). But this means that as soon as we try to generalize over “c” we must address a preliminary issue glossed over in the argument-sketch: Is the skeptic going to use a single skeptical hypothesis, such that each perceptual belief entails the denial of this hypothesis? Or is he going to use many different skeptical hypotheses, each one tailored to suit the specific perceptual belief at hand? In other words, the skeptic might be thinking along the lines of

> There is a skeptical hypothesis that makes trouble for every perceptual belief;

or he might be thinking along the lines of

> For every perceptual belief there is some trouble-making skeptical hypothesis.

Since the skeptic might try either option, we’ll have to distinguish between two types of skeptical strategy, depending on what sort of hypotheses the skeptic wants to employ.

Let’s first consider the strategy of the *general skeptical hypothesis*. A skeptic pursuing this strategy wants to advance a single skeptical hypothesis such that each perceptual belief, each belief in a proposition of class C, entails the denial of this hypothesis. Interestingly, it is not entirely straightforward to extract a skeptical hypothesis from a skeptical scenario that will do this job. The following is a somewhat rough proposal:

H: There is an evil demon who makes it the case that human agents tend to hold beliefs committing them to the view that there are material objects even though there are no material objects.

The idea is that our perceptual beliefs all entail the proposition that there are material objects, which of course entails the denial of H.[14]

The proposal is a bit rough. There may well be perceptual beliefs that do not entail the existence of material objects. For example, various conditional beliefs and/or negative beliefs might depend epistemically on how things perceptually appear to us, thereby counting as perceptual beliefs, but might not themselves entail the existence of any material object, even if they are based on beliefs that do. If so, a skeptic employing hypothesis H will be forced to somehow restrict C, the class of propositions he is targeting, which would of course weaken the skeptical conclusion. (One might note in passing that propositions that explicitly mention material objects, say the proposition that there are chairs in the room, entail the proposition that there are material objects only under the assumption that material objects (chairs) are necessarily material objects: though this seems eminently plausible, it might still be noteworthy that a skeptic employing H relies on this assumption.) Rather than going into other worries one might possibly have about hypothesis H, we turn to the skeptical argument that employs H.[15]

A more carefully worked out skeptical argument that (a) is based on Closure 6, (b) proceeds along the lines indicated by the argument-sketch given above, (c) employs a single general skeptical hypothesis like H, and (d) aims at the general skeptical conclusion saying that human agents don't have any perceptual knowledge (don't know the propositions of class C), will go like this:

Argument A:

P1. $\forall S, \forall p, \forall q$: SKp and SK(p entails q) and SBq bod from p and (p entails q) $\rightarrow$ SKq
P2. $\forall S$: ~SK~H
3. $\forall S, \forall p$: ~SKp or ~SK(p entails ~H) or ~(SB~H bod from p and (p entails ~H))
P4. $\forall S, \forall p \in$C: SBp $\rightarrow$ SK(p entails ~H)
P5. $\forall S, \forall p \in$C: SBp $\rightarrow$ SB~H bod from p and (p entails ~H)
6. $\forall S, \forall p \in$C: SBp $\rightarrow$ ~SKp
7. $\forall S, \forall p \in$C: ~SKp

Argument A is clearly valid. If it were sound, it would establish the strong conclusion the skeptic is aiming for. But the argument is not sound. On the contrary, premises P4 and P5 are ridiculously false.

Consider premise P4. Is it plausible to suggest that human agents know of each one of their perceptual beliefs that it entails the denial of the skeptical hypothesis? This ascribes all sorts of entailment beliefs to agents (who hold perceptual beliefs), and it is perfectly clear that typical agents do not have all these entailment beliefs. It is still more clear that it is not the case that all agents have all these entailment beliefs about the perceptual propositions they believe—it is even doubtful that most of the agents who happen to be aware of the skeptical hypothesis (i.e. philosophers) have a large number of such entailment beliefs. Consider premise P5. It says that all agents (who hold perceptual beliefs) believe the denial of the skeptical hypothesis, ~H, and believe it based on deduction from each of their premise-belief pairs *p* and (*p* entails ~H), for every perceptual proposition *p* they believe. This ascribes all the entailment beliefs already ascribed by premise P4 to all agents. In addition, it ascribes beliefs in the denial of the skeptical hypothesis to all agents holding any perceptual beliefs. In addition, it maintains that each agent's belief in the denial of the skeptical hypothesis is radically overdetermined, being based on each of the premise-belief pairs *p* and (*p* entails ~H), for each of the agent's perceptual beliefs *p*. Is any of this plausible? The answer is obvious: It is not at all plausible.

Let us look at the version of Argument A that takes the slightly safer principle Closure 7 as its starting point. The only place where this will make a difference is premise P5—the difference being that Closure 7 will make things worse. The Closure 7 version of P5 requires that, for *each* of their perceptual beliefs *p*, agents believe the denial of the skeptical hypothesis, ~H, *solely* on the basis of deduction from the premise-pairs *p* and (*p* entails ~H). For any agent holding more than a single perceptual belief *p*, this isn't even logically possible.

What about the version of Argument A that starts with Closure 6.5? Again, this can only make a difference at premise P5. The Closure 6.5 version of P5 is not as absurd as the Closure 7 version. It is however still worse than the Closure 6 version, for it requires that agents believe ~H based on deduction from the premise-pairs *p* and (*p* entails ~H), for each of the agents' perceptual beliefs *p*, and that they don't also believe ~H based on any defective reason.

Let us now consider the strategy of the *particular skeptical hypotheses*. Judging from the literature on closure principles and their role in skeptical arguments, this strategy appears to be the preferred one. Authors who discuss this topic typically make use of some argument-sketch, albeit one based on Closure 2 rather than Closure 6, 6.5, or 7. They select some perceptual proposition, say the proposition that there are chairs in the room, supposed to be believed by some agent. They then extract from the skeptical scenario a skeptical hypothesis tailored to suit the particular belief under consideration; e.g. an hypothesis to the effect that the evil demon through his machinations makes it look to the agent as if there are chairs in the room even though there aren't.[16]

The recipe for constructing such a particular skeptical hypothesis is this. Take a particular agent, S, and a particular proposition, c, say the proposition that there are chairs in the room. Assume c belongs to the skeptic's target class, C, and assume S believes c. The particular skeptical hypothesis corresponding to c for S is:

H⟨S,c⟩: There is an evil demon who misleads S into believing c even though c is false.

Hypothesis H⟨S,c⟩ is "the particular skeptical hypothesis corresponding to c for S" because it applies to particular S and because proposition c itself occurs within the formulation of the hypothesis—actually, the hypothesis says explicitly that c is false. Note that, in virtue of this feature, our proposition c trivially entails the denial of H⟨S,c⟩. But note also that most other propositions from the target class C, propositions that are logically independent from c, will not entail the denial of this hypothesis. To reach a general skeptical conclusion, one covering all beliefs from C, the skeptic will need a different particular hypothesis, H⟨*S*,*p*⟩ for each proposition *p* in the target class C. Thus, a skeptic who pursues this strategy and wants to give a general argument, based on Closure 6, for the conclusion that we don't have any perceptual knowledge will have to make use of

Argument B:

P1. ∀*S*,∀*p*,∀*q*: *S*K*p* and *S*K(*p* entails *q*) and *S*B*q* bod from *p* and (*p* entails *q*) → *S*K*q*
P2.1. ∀*S*,∀*p*∈C: ~*S*K~H⟨*S*,*p*⟩
3.1 ∀*S*,∀*p*∈C: ~*S*K*p* or ~*S*K(*p* entails ~H⟨*S*,*p*⟩) or ~(*S*B~H⟨*S*,*p*⟩ bod from *p* and (*p* entails ~H⟨*S*,*p*⟩)
P4.1. ∀*S*,∀*p*∈C: *S*B*p* → *S*K(*p* entails ~H⟨*S*,*p*⟩)
P5.1 ∀*S*,∀*p*∈C: *S*B*p* → *S*B~H⟨*S*,*p*⟩ bod from *p* and (*p* entails ~ H⟨*S*,*p*⟩)
6. ∀*S*,∀*p*∈C: *S*B*p* → ~*S*K*p*
7. ∀*S*,∀*p*∈C: ~*S*K*p*

Argument B is a bit more complicated than Argument A and even more implausible. Its crucial premises, P4.1 and P5.1, assume, quite absurdly, that agents believe all the denials of the particular skeptical hypotheses corresponding to their perceptual beliefs and believe all the entailments from their perceptual beliefs to the corresponding denials of the particular skeptical hypotheses. That is, they ascribe to agents lots of beliefs involving denials of lots of different particular skeptical hypotheses, which is even more implausible than the beliefs ascribed to agents by the corresponding premises of Argument A.

The Closure 6.5 version of premise P5.1 again makes things a bit worse by requiring in addition that agents must not have any defective reasons for their beliefs in all the denials of the particular skeptical hypotheses.

The only point where the strategy of the particular skeptical hypotheses makes a small positive difference compared to the strategy of the general skeptical hypothesis is with respect to the Closure 7 based version of the argument. Closure 7's version of P5.1 will say that for each of their perceptual beliefs p, agents S believe $\sim H\langle S,p\rangle$, and that each of these beliefs in $\sim H\langle S,p\rangle$ is based *solely* on one pair of premises—a different pair for each denial. At least this is not impossible, although it is of course still quite absurd.

One might object that a skeptic who pursues the strategy of the particular skeptical hypotheses is not committed to the view that there is a single skeptical argument, like Argument B, establishing in one stroke that all the perceptual beliefs of all agents fail to constitute knowledge; instead, the skeptic is merely committed to the view that for each perceptual belief of each agent there is some skeptical argument establishing, by way of a particular skeptical hypothesis, that the agent doesn't know the proposition corresponding to the hypothesis. Fair enough, but it will not make a difference: most of the arguments envisaged by this response will be unsound, ascribing to agents beliefs they do not hold.

We pause for a moment to reconsider a point about the format used here to represent closure-based skeptical arguments. The point questions the adequacy of representing perceptual beliefs in terms of a skeptical target class C of perceptual propositions p believed by agents S. One might object (and Robert Audi has so objected) that this is inadequate because the adjective "perceptual" is not plausibly construed as picking out a property of propositions; instead, it modifies the relation of *believing* which holds between agents and propositions. In other words, there are no perceptual propositions as such: pretty much any proposition, whatever its nature or content, can be believed by an agent on perceptual grounds (viz. by testimony).[17]

At first it looks like accommodating this objection should require only relative mechanical changes in Arguments A and B: replace all occurrences of the restricted quantifier "$\forall p \in C$" by unrestricted "$\forall p$"; and replace all occurrences of "$S\mathrm{B}p$" by "$S\mathrm{PB}p$", saying that S perceptually believes p. (Both arguments would then directly proceed to the conclusion "$\forall S, \forall p$: $S\mathrm{PB}p \rightarrow \sim S\mathrm{K}p$", saying that all agents fail to know any proposition they perceptually believe.) On second thought, however, accommodating the objection presents a bit more of a problem—a problem for the skeptic, that is, in coming up with a skeptical hypothesis. Take the skeptic who wants to employ a single general skeptical hypothesis. If pretty much any proposition can be believed on perceptual grounds, including necessary truths, the skeptic must produce a skeptical hypothesis whose denial is entailed by any necessary truths that happen to be believed by sundry agents on perceptual grounds (viz. by testimony). No hypothesis along the lines of H will do this job. Moreover, any candidate hypothesis must immediately face the objection that it is not a possible hypothesis since its denial, being entailed by a necessary truth, must be a necessary falsehood. The difficulty remains essentially the same for a skeptic who wants to employ different skeptical hypotheses, each one tailored to suit a particular proposition believed by an agent on perceptual grounds, including any necessary truths believed on perceptual grounds. We leave this difficulty unresolved and return to the main line of the paper.

We have seen that Closures 6 to 7, however plausible they might be as closure principles, will not take the skeptic to the general conclusion that no one has any perceptual knowledge via worked out versions of the standard sketch of closure-based arguments, i.e., via arguments like A or B. The attempt to strengthen the antecedent of some relatively simple closure principle in order to handle the belief and warrant problems just forces the skeptic to make a number of additional assumptions when trying to reach the skeptical conclusion $\sim S\mathrm{K}p$ in its general form.

The additional assumptions are rather difficult to defend.

The problem we have been belaboring for the last few pages, let us call it the "generalization problem", will always arise to various degrees for arguments proceeding along the lines considered so far—with the sole exception being arguments based on the clearly flawed principle Closure 1. All other closure principles strengthen Closure 1's antecedent in various ways, thereby forcing the skeptic to ascribe problematic beliefs to agents. The generalization problem arises with respect to all these principles and the rough rule seems to be: the safer the principle the more severe the problem.

We note that the generalization problem, though somewhat neglected, is a rather obvious one. Part of the reason why it is not given the attention we think it deserves might be this. While it is usually acknowledged that Closure 1 is not an acceptable principle (the preferred candidate usually being the slightly more plausible Closure 2), there seems to be a tendency to forget about this point and to inadvertently revert back to Closure 1 when it comes to thinking about the skeptical argument that is supposed to be based on "the" closure principle.

So far we have been arguing that closure principles with a strengthened antecedent will not allow the skeptic to reach the general conclusion that we don't have any perceptual knowledge by way of a closure-based argument like the ones considered above. However, this does not quite show that such principles are of no use at all to the skeptic. Maybe an argument based on a closure principle can be useful to the skeptic by taking him part of the way, where the other part must then be covered by some additional consideration.

To get a better picture of the two-stage strategy we have in mind here, turn back to the Argument Sketch based on Closure 6 given earlier. Instead of interpreting it as a sketch of an argument by arbitrary instance, which involves ascribing to agents lots of beliefs about skeptical hypotheses they do not hold, why not apply it to just a few (non-arbitrary) agents who are aware of skeptical hypotheses? Or even better, apply it just to one agent, say, to *me*—after all, I am aware of skeptical hypotheses. One of the perceptual propositions I believe is that there are chairs in the room. Applying the argument just to myself and this particular proposition, so the idea would go, I must conclude that I don't know that there are chairs in the room; and now comes the additional step: "If I don't know *that*, then I don't know any perceptual proposition; and if *I* don't know any perceptual, then no one knows any perceptual proposition".

Before considering the crucial additional step which constitutes the second stage of the overall argument, one should note that the first stage is already a bit difficult. Yes, I am aware of the skeptical hypothesis, h, be it the general hypothesis or the particular hypothesis corresponding to the proposition that there are chairs in the room. I believe that this proposition, call it c, entails ~h and, let's assume, I know that it does. Thus, premise P4 of the closure-based argument applied to me and my perceptual belief c seems fairly safe. But what about premise P5? It says that I believe ~h based, at least in part, on deduction from my belief c together with my belief that c entails ~h. Do I? I find this rather difficult to tell. If it is supposed to mean that I wouldn't believe ~h if I didn't believe c and (c entails ~h), or that the strength of my belief in ~h would then be less than it is now, it feels wrong. But again: it is hard to tell. And it does not get any easier when I ask myself, à la Closure 6.5, whether I have any defective reasons for believing ~h. One thing that is very clear to me, though, is that I don't believe ~h based *solely* on c and (c entails ~h); so Closure 7 is out of the running. Assuming that you are roughly like me, we can observe that the first shot at the new, two-stage strategy tends to bog down at the first stage.

But maybe one can cut through this problem by reshaping the original idea behind the two-stage strategy. Turn to one of the general closure-based arguments, say Argument A. It does not establish the general skeptical conclusion that no agent has any perceptual knowledge. But it is not implausible to maintain that it establishes a still general but weaker, *conditional* conclusion:

CA: *If* SK(p entails ~H) and SB~H bod from p and (p entails ~H), *then* ~SKp,

which, note, is to be taken as a general principle applying to all agents, S, and all propositions, p, such that $p \in$C and SBp.

Evidently, for this principle to hold it does not matter whether there are any agents who actually believe that p entails ~H, or believe ~H based on deduction from p and (p entails ~H). The principle circumnavigates these issues by conditionalizing on the problematic premises P4 and P5 of Argument A. The principle does not, of course, conditionalize on the other two premises of Argument A. So it depends on the relevant closure principle, that is, on Closure 6 (though analogous principles depending on Closure 6.5 or Closure 7 could easily be formulated), and it depends on the skeptic's contention that no one knows ~H. (The modal status of CA is somewhat unclear. Since Argument A is valid, the principle could safely be declared to be necessary, if the two premises it depends on were clearly necessary. The relevant closure principle, let's assume, is necessarily true if true at all. But it is not clear whether the same can be said for the premise that no one knows ~H; it is not even clear whether the skeptic typically takes this premise to be a necessary truth.)

CA is of course not the strong unconditional conclusion the skeptic is aiming for. It does not say that no one has any perceptual knowledge. All it says is, roughly speaking, that agents who satisfy CA's antecedent don't know the perceptual propositions that satisfy CA's antecedent.[18] The skeptic needs more than this. And he might maintain that he can get it. With CA in hand, he might want to proceed to the second stage of his argument by making the following additional claim:

Claim: Principle CA shows that (actual or possible) agents who satisfy the antecedent of CA don't know any perceptual propositions that satisfy the antecedent of CA; but if *such* agents don't know *these* perceptual propositions, then no agent knows any perceptual proposition.

An observation may help to clarify the status of this Claim. Even if the skeptic could not get anything more than the conditional conclusion CA, this would not by itself amount to an argument *against* skepticism. That is, even if the skeptic could show nothing more than that agents who satisfy CA's antecedent fail to know the relevant propositions, that would not show that agents who don't satisfy CA's antecedent do know the relevant propositions, (that the skeptic can only show ~SKp for agents satisfying the antecedent doesn't show that we can establish that *only* agents satisfying the antecedents are such that ~SKp). This much is surely right and should be kept in mind by anyone who wishes to resist closure-based skepticism. However, the skeptic's Claim does much more than just making this observation. It says in effect that, because agents who satisfy the antecedent do not know the perceptual propositions that satisfy the antecedent, no one knows any perceptual proposition.

So, what about this Claim? It is a curious one. Note that we would not normally be inclined to reason thus: "Everything that is both F and G is not K; so, Everything is not K." The skeptic's Claim, however, does seem to embrace this sort of reasoning. Of course, with the help of a further premise, a premise of the form "Everything that is *not* both F and G is not K either", the desired result would be immediately forthcoming. But no such further premise is available to the closure-based skeptic under consideration here.[19] So to defend his Claim, the skeptic would have to point to a special feature—some special property of CA—that makes this sort of reasoning acceptable in this special case. What feature could that be? Well, if the skeptic could make a case for the idea that (actual or possible) agents who satisfy the antecedent of CA are somehow ideally situated, epistemically speaking, with respect to the relevant perceptual propositions, i.e., with respect to the perceptual propositions from which they have deduced the denial of the skeptical hypothesis, then he might be able to conclude that, if such agents don't know these propositions, no one knows any perceptual propositions. But it is hard to see how such a case should be made. On the face of it, there is nothing in the antecedent of CA suggesting that agents satisfying it are in an especially good epistemic position with respect to the relevant propositions.

(One might even try to argue the contrary making use of the idea that knowledge is vulnerable to *defeat*: i.e., a person who knows *p* can lose that knowledge upon acquiring a defeater for the reasons or grounds on which that knowledge was based. It is not very clear to anyone what "acquiring a defeater" amounts to. But it might be held that to acquire a defeater it is sufficient, at least under certain circumstances, that the person has considered it—toyed with it mentally—even if she doesn't believe it, or even disbelieves it. If so, one could respond to the skeptic by arguing that agents who satisfy the antecedent of CA are agents who are in an epistemically *bad* position with respect to *p*. They have toyed with H and, even thought they disbelieve H, they have thereby acquired a defeater for *p*: you know less by reflecting more.)

We won't go any further into this issue. We are satisfied if we have managed to show that even the best knowledge-closure principles (of the "strengthened antecedent" type) are of more limited use to the skeptic than one might have thought initially. Even assuming that Closures 6, 6.5, or 7 are plausible principles—and of course always granting that we do not know skeptical hypotheses to be false—the best the skeptic can do based on these principles alone is to get to a conditional conclusion such as CA. To get any further, the skeptic needs to support something like his additional Claim, and this seems to require a new argument, an argument we haven't seen yet.

Before we leave this section, we want to return, once more, to general closure-based arguments derived from the Argument Sketch given above and discuss a skeptical strategy that is not quite covered by Arguments A or B. Remember the general skeptical hypothesis

H: There is an evil demon who makes it the case that human agents tend to hold beliefs committing them to the view that there are material objects even though there are no material objects.

Earlier we looked at an argument, Argument A, which requires the assumption—an assumption that is not entirely trivial—that each of our perceptual beliefs entails the denial of H. One may have noticed that there are other propositions entailing the denial of H. In particular, there is the general proposition

X: There are material objects.

This proposition has the advantage that it trivially entails the denial of H. A skeptic might want to make use of this proposition in a Closure 6 (or 6.5 or 7) based argument. To do so, he needs an argument different from Argument A or B. For Closure 6, the new argument would go like this:

Argument C:

P1. $\forall S, \forall p, \forall q$: SKp and SK(p entails q) and (SBq bod from p and (p entails q)) $\rightarrow$ SKq
P2. $\forall S$: ~SK~H
3. $\forall S, \forall p$: ~SKp or ~SK(p entails ~H) or ~(SB~H bod from p and (p entails ~H))
P4.2 $\forall S$: SBX $\rightarrow$ SK(X entails ~H)
P5.2 $\forall S$: SBX $\rightarrow$ SB~H bod from X and (X entails ~H)
6.2 $\forall S$: SBX $\rightarrow$ ~SKX
7.2 $\forall S$: ~SKX

Argument C concludes that no one knows the proposition that there are material objects.[20] It does better than Arguments A and B because it ascribes fewer beliefs to agents. Indeed, it seems that there are agents, mostly philosophers, who do hold the beliefs ascribed by premises P4.2 and P5.2. However, the argument is still unsound because it ascribes these beliefs to all agents who believe that there are material objects, not just to philosophers. Moreover, there is also the worry whether agents who do believe X, and (X entails ~H), and believe ~H really believe the latter based on deduction from the former, as maintained by P5.2.

Of course, the conclusion of Argument C, that no one knows that there are material objects, though an interesting skeptical conclusion in its own right, is not the conclusion the skeptic is ultimately aiming for; it is not the conclusion that no one has any perceptual knowledge. What is the skeptic's plan for reaching that conclusion?

The most obvious plan that comes to mind would be to argue that no one knows any perceptual proposition because no one knows X. That is, the most obvious plan would be to employ Closure 6 again, this time in argument modeled on A, but with the second premise of Argument A replaced by the conclusion of argument C:

P1. $\forall S, \forall p, \forall q$: SKp and SK(p entails q) and SBq bod from p and (p entails q) $\rightarrow$ SKq
7.2 $\forall S$: ~SKX
3. $\forall S, \forall p$: ~SKp or ~SK(p entails X) or ~(SBX bod from p and (p entails X))
P4.3 $\forall S, \forall p \in$C: SBp $\rightarrow$ SK(p entails X)
P5.3 $\forall S, \forall p \in$C: SBp $\rightarrow$ SBX bod from p and (p entails X)
6. $\forall S, \forall p \in$C: SBp $\rightarrow$ ~SKp
7. $\forall S, \forall p \in$C: ~SKp

Premises P4 and P5 of this argument are less implausible than the corresponding premises in the original version of A. Surely, it is more likely that agents holding perceptual beliefs also hold the relevant beliefs involving X than that they hold the relevant beliefs involving ~H. Note, however, that it is still quite implausible to maintain that such agents believe X based on deduction from *p* and (*p* entails X), for each of their perceptual beliefs *p*. Moreover, the resulting overall argument—Argument C combined with this continuation—does not appear significantly more plausible than Argument A.

An intriguing alternative to this way of continuing Argument C might be to continue not with an argument based on Closure 6, but with an argument based on a closure principle that goes all the way back to Closure 1; more precisely, to an alternative way of strengthening the antecedent of Closure 1 that we have not considered yet:

Closure 1A: (K*p* and (*p* analytically entails *q*)) $\rightarrow$ K*q*.

For those who still believe in analyticity, this will be a good candidate for a plausible closure principle, provided they are prepared to construe "analytic entailment" in a certain manner. First, the notion of analyticity must make sense as applied to propositions, and not just to linguistic entities like sentences. Second, the relation of analytic entailment must be taken to imply that one cannot possibly believe the analytically entailing proposition without believing the entailed proposition. Third, the relation of analytic entailment must be taken to imply that one cannot possibly know the analytically entailing proposition without knowing the entailed proposition.

If there is analytic entailment, and if it can be construed in this way, then Closure 1A does not raise the belief and warrant problems for the range of propositions to which it applies. What is this range? There are the seemingly trivial, "decompositional" candidates, e.g.: the proposition that Audumla is a brown cow analytically entails the proposition that Audumla is a cow.[21] Then there are the "standard" candidates, e.g.: the proposition that Leopold is a bachelor analytically entails the proposition that Leopold is unmarried. But the sort of cases needed by the skeptic are a bit different still. The skeptic who wants to continue Argument C with a closure argument based on Closure 1A must maintain that all perceptual propositions, such as the proposition that there are chairs in the room, analytically entail the proposition that there are material objects. These candidates for analytic entailment are not at all like the brown cow/cow cases, nor are they quite like the bachelor/unmarried cases; they are cases of "philosophical analyticity".[22]

Of course, *post* Quine, pretty much everything about analyticity is contentious, with candidates of philosophical analyticity surely being the most contentious. Note also that even friends of analyticity might well oppose the idea that analytic entailment (especially the philosophical variety) can be construed in the manner required, namely as implying that certain propositions cannot possibly be believed without believing certain other propositions and cannot possibly be known without knowing these other propositions.[23] Moreover, a skeptic who is prepared to take on the anti-analyticity camp and to defend the required construal of analytic entailment still faces the problem that the first leg of his overall argument, Argument C, appears to be unsound.

Closure 1A is perhaps the "leanest" closure principle with strengthened antecedent whose skeptical potential is worth exploring. Let us mention, for comparison, what would be the "fullest" closure principle of this sort. It would be a principle taking care of any challenges one might come up with even for Closures 6.5 or 7 by restricting the antecedent to agents who believe *q* solely based on deduction from *p* and (*p* entails *q*) and for whom "nothing goes wrong" (epistemically speaking) in so believing *q*:

Closure 8: (K*p* and K(*p* entails *q*) and B*q* solely based on deduction from *p* and (*p* entails *q*) and "nothing goes wrong") $\rightarrow$ K*q*.

This closure principle should be the safest of the "strengthened antecedent" family. Of course, the safety is paid for by the considerable (excessive?) vagueness of the "nothing goes wrong"-

clause which gestures towards potential challenges to Closures 6.5 and 7 and stipulates that none of them apply.

It is clear that Closure 8 will not be of use in a skeptical argument in the style of Argument A or B. Such arguments will have the special problems noted for arguments based on Closure 7 (concerning the "solely based on" clause) and they will suffer from an exacerbated form of the problems noted for arguments based on Closure 6.5; that is, they will require a premise to the effect that nothing at all goes wrong (epistemically speaking) with agents who believe *q* based on…. Of course, a conditional principle might be derived, analogous to principle CA but incorporating all elements of the antecedent of Closure 8. This would then, again, require supplementation along the lines of the skeptical Claim mentioned earlier.

Our main question, we said, is this: Are there any plausibly true epistemic closure principles concerning knowledge that figure centrally in arguably sound skeptical arguments with interestingly strong conclusions? What is our answer so far? Well, we have found some knowledge-closure principles (of the strengthened-antecedent variety) that are fairly plausible. But no such principle, we have argued, will by itself take the skeptic to the strong conclusion that no one has any perceptual knowledge: as far as knowledge-closure principles of this sort are concerned, there is, one might put it, no skeptical argument that is both sound and "pure". The skeptic would need to supplement his argument with an additional Claim the status of which is unresolved.

IV. Plausible Principles (b): Weakening the Consequent

The closure principles discussed in Section III all try to avoid the belief problem (and the warrant problem) by strengthening the antecedent of Closure 1 in various ways. We now consider the second major strategy for avoiding the belief problem facing knowledge-closure principles: weakening the consequent. All principles discussed in this section respond to the belief problem by weakening the consequent of the principle so that it does not imply belief. The more plausible principles in this section will respond to the warrant problem with various strengthenings (compared to Closure 1) or reinterpretations of the antecedent of the principle.

Let K* stand for a relation weaker than knowing: let it stand for the relation of "being positioned to know" or "being in a position to know" where these locutions are understood not to imply belief in the proposition in question. Being "positioned to know" that Leopold is in the room does not imply that one believes that Leopold is in the room. Perhaps a closure principle making use of this K* locution will be of use to the knowledge-skeptic in formulating a plausible closure principle that is of possible use in skeptical argumentation.

We saw in the previous section that closure principles with cluttered antecedents are unlikely to be of use in skeptical argumentation. We therefore begin this discussion with the leanest relevant closure principle making use of the K* locution:

Closure 9 : $(Kp \text{ and } (p \text{ entails } q)) \rightarrow K^*q$

This says that (necessarily) all agents are positioned to know propositions entailed by propositions they know.[24] With its lean antecedent, this principle, if plausible, would be well suited for serving in a skeptical argument for the general conclusion that no one has any perceptual knowledge. Two immediate issues arise, however.

The first issue arising as we begin discussion of "weaken the consequent" closure principles is this. The point of introducing the K* terminology is so that the consequent of a closure principle using this notion is strictly *weaker* than the consequent of a closure principle with the standard Kq consequent. Because of this, the key skeptical claim that commences a skeptical argument (before ~K~SH; now ~K*SH) must be strictly *stronger* when using a closure principle of this new type. Where before the skeptical argument began with the claim that we do not know the denial of the skeptical hypothesis, the argument must now begin with the claim that we are not even positioned to know this denial. Anti-skeptics tempted to respond to the skeptic by claiming that we do know the denial of the skeptical hypothesis will surely not be able to resist arguing that we are at least positioned to know this denial. If the skeptical claim becomes very strong (because, eg, the interpretation of "positioned to know" makes the consequent of the closure principle exceptionally weak) it may well become inevitable that the anti-skeptic will want to make this move. For now, however, we generously set this issue aside.[25]

The second and perhaps most obvious issue arising because of the introduction of this "positioned to know" terminology is this. Serious evaluations of principles involving the K* locution and arguments employing these principles will depend upon how "positioned to know" is understood. We will discuss several readings of this phrase in this section. We will argue that many natural readings of "positioned to know" are clearly of no use in skeptical argumentation. Additionally we will argue that some natural readings lead to serious interpretive and internal difficulties with the notion of "positioned to know" itself. Though we will identify a few not wholly implausible readings of "positioned to know" and will find at least one skeptical argument worth discussion, we will find nothing promising for the skeptic in this section.

First, however, we announce a partial shift in methodology. Recall that near the beginning of this paper we chose a strategy of looking for plausible closure principles and then evaluating their possible use in skeptical argumentation. Because of our progress in Section III, we are now familiar with several problems closure-based skeptical arguments can encounter. It seems wise to apply this knowledge now by shifting somewhat towards a strategy of looking for a plausible closure-based skeptical argument and then working backwards to see if the closure principle needed for such an argument is plausible. This shift began when we observed above that the antecedent of Closure 9 makes it a good candidate for overcoming some of the concerns raised in the previous section's discussion of strengthened antecedent closure principles. We won't immediately begin looking for a skeptical argument. Rather, we'll first explore a few general issues involving "positioned to know" and closure principles employing "K*". After doing this we will move on fairly quickly to formulate the most knowledge-skeptic friendly closure principles employing K*.

In order to evaluate Closure 9 and other weakened consequent closure principles we must at least begin to fill in some content for the K* locution. Understandably, we encounter some difficulty at this point in our inquiry. There is no standard analysis of what it is to be "positioned to know" something that we can plug into the principle. Additionally, it seems clear that the phrase "being positioned to know" is in many ways vague and slippery. Perhaps we can safely stipulate the following two basic components involved in the understanding of this locution. First, if one knows something then one is positioned to know it (knowing is the limit case of being positioned to know – one is especially well positioned to know what one knows). Second, one is positioned to know p only if p is true.[26] Beyond these two fairly minor points, it is not immediately obvious even what sort of analysis of "positioned to know" is most worth exploring, let alone what specific analysis merits attention.

The weakening the consequent strategy is introduced as a possible way to avoid the belief problem that plagues standard formulations of knowledge-closure principles. Since this problem concerned the possible absence of a belief in *q* in earlier closure principles, perhaps it is natural to begin with an "all but belief" interpretation of K*. Perhaps one is positioned to know something if one has (at least) everything required for knowing it except belief; i.e., K* = K minus B.[27] This reading has this advantage: because on this reading "K*p" is differentiated from "Kp" only by the psychological issue of whether belief is present, the key skeptical claim (viz., ~K*~SH) is on this reading only barely stronger than in the strengthened antecedent arguments (~K~SH). This natural first reading of K* that is somewhat attractive.

Given this natural interpretation of K*, Closure 9 suffers from at least two difficulties. First, just as Closure 1 made it too easy to know necessary truths, Closure 9 appears to make it too easy to be positioned to know necessary truths. If Closure 9 is true, anyone who knows anything will be thereby positioned to know all necessary truths. Second, and more generally, Closure 9 seems to make it too easy to be positioned to know things. It doesn't seem plausible that an agent who knows p is thereby positioned to know any entailed proposition q because this would imply that the agent is positioned to know q even if the agent doesn't even believe that q is so entailed.

These reflections lead, unsurprisingly, to a weakened consequent principle modeled on Closure 2:

Closure 10: Kp and K(p entails q)) $\rightarrow$ K*q

This is probably the sort of principle that most advocates of the "weakening the consequent" strategy have in mind.[28] On the face of it, Closure 10 looks like at least a superficially plausible closure principle. Whatever its merits *as a closure principle*, however, the lessons of the Section III teach us that this principle will not be of use in a skeptical argument. Note that in addition to its weakened consequent this principle involves a strengthened antecedent. A skeptic attempting to employ this principle in a general skeptical argument will be caught ascribing to agents all sorts of entailment beliefs involving the denial of skeptical hypotheses that typical agents are unlikely to hold. This is an important reminder that those looking for a knowledge-closure principle of possible use in a skeptical argument must strive to keep the antecedent of the closure principle as uncluttered as possible.

Another serious difficulty for those interested in using principles like Closure 9 and Closure 10 for skeptical purposes should be noted. Given the somewhat rough "all but belief" interpretation we have suggested for K*, the most natural way to understand K* is as a disguised conditional. "K*p" on this reading seems equivalent to (necessarily) [Bp $\rightarrow$ Kp]. "S is positioned to know p" means, on this reading, "S is such that, necessarily, if S believes p then S knows p".[29] Interpreting K* in this way is not a promising strategy for the skeptic to pursue. Depending on the strength of the conditional, K*-involving closure principles are be equivalent to or imply principles identical to or relevantly like those we have already seen and rejected.[30] For example, reading the conditional as a strict conditional,

Closure 9: (Kp and (p entails q)) $\rightarrow$ K*q
is equivalent to
Closure 9': (Kp and (p entails q)) $\rightarrow$ (Bq $\rightarrow$ Kq)
which implies
Closure 9'': (Kp and (p entails q) and Bq) $\rightarrow$ Kq

Closure 9'' is an implausible "strengthened antecedent" principle with an antecedent weaker than the uninspiring Closure 5 from Section III. Because even the stronger antecedent Closure 5 is false, Closure 9'' must be false which implies that Closure 9 is false and therefore useless to the skeptic. Similarly, the more plausible looking Closure 10, on this "all but belief" conditional reading of K* straightforwardly implies

Closure 5: (K*p* and K(*p* entails *q*) and B*q*) $\rightarrow$ K*q*.

We saw in Section III that Closure 5 is false. It follows that any principle implying Closure 5 is also false. It seems that the skeptic will have to look beyond the natural "all but belief" conditional reading of K* in attempting to formulate a skeptical argument employing K*. We will return to this issue very soon as we get closer to finding a principle that can overcome the other barriers we seem to have in front of the skeptic.

First let's take note of the problems Closure Principles 9 and 10 have because of their different commitments concerning an agent's epistemic status relative to the entailment of q by p. Closure 9 seemed quite problematic as a principle because it says *nothing* about one's epistemic status relative to this entailment relation. Closure 10, which specifies that S knows the relevant entailment relation, seemed more promising as a principle, but is of no use in a skeptical argument at least because it implies universal belief in the specified entailment relation. It seems that any promising skeptical strategy must lay intuitively "in between" Closure 9 and Closure 10 regarding S's epistemic standing regarding the entailment relation. It would seem natural for the skeptic to look at this principle which makes use of the K* device in both its consequent and its antecedent:

Closure 11: (Kp and K*(p entails q)) $\rightarrow$ K*q

Closure 11 says that one who knows p and is (at least) positioned to knows (p entails q) is thereby (at least) positioned to know q.[31] This principle avoids the belief-entailing problematic features of Closure 10 but simultaneously requires that one stand in a positive epistemic relation to the entailment. This should help the skeptic avoid attributing problematic entailment beliefs while also at least partially addressing the worry that the principle might make it too easy to count as positioned to know things. Closure 11 might on these grounds be thought to be the most promising principle we have seen so far.

Closure 11 is, from a structural point of view, quite promising when looking ahead to the issue of general skeptical argumentation. But one thinking that Closure 11 is promising on the current interpretation of K* is mistaken. First a minor problem: Closure 11 may well be too close to Closure 9 and might therefore inherit some of the difficulties facing Closure 9. Specifically, Closure 11, like Closure 9, may make it too easy to count as positioned to know something. Consider an agent who knows p and is *merely* positioned to know that p entails q (but who doesn't even believe, let alone know, that this entailment holds). Given Closure 11, this agent counts as positioned to know q, but this seems too generous. It seems that rather than being positioned to know q, our agent is at most positioned to be positioned to know q (whatever exactly that might mean). We will see one possible response to this worry later when we consider a different reading of K*. First, however, we turn to a larger problem for Closure 11 on the current reading.

Recall that Closure 10 on the current reading of K* implies the implausible Closure 5:

Closure 10: Kp and $K(p$ entails $q)) \rightarrow K^*q$
Closure 5: ($K\textit{p}$ and $K(p$ entails $q)$ and $Bq) \rightarrow Kq$.

Closure 10 must be rejected because it implies the implausible Closure 5.
Closure 10, however, is a strengthened antecedent version of Closure 11:

Closure 11: (Kp and $K^*(p$ entails $q)) \rightarrow K^*q$.

This implies that if Closure 10 is false then so is Closure 11. Closure 11 must therefore be rejected.

Strictly speaking, of course, what we reject in rejecting Closure 11 is Closure 11 with K* interpreted in the "all but belief" manner we have to this point adopted (on any reading of the relevant conditional). It seems that one looking for a plausible weakened consequent closure principle for possible use in skeptical argumentation must do one or both of the following:
(a) Choose a different interpretation of K*,
(b) Choose a closure principle structurally different from Closure 11.
With an eye on skeptical argumentation, we are confident that best move for the skeptic at this point is to choose option (a) but reject option (b).

Structurally, after all, Closure 11 has some nice features that tell against choosing option (b). Indeed, from the point of view of the structure of the overall skeptical argument, Closure 11 has all the features the skeptic is looking for. The sketch of a skeptical argument using Closure 11 is familiar and in some ways promising. Using the same abbreviations as before, the sketch of a skeptical argument using Closure 11 looks like this:

P1. (Kp and $K^*(p$ entails $q)) \rightarrow K^*q$
P2. ~K*~h
3. ~Kc or ~K*(c entails ~h)
P4. K*(c entails ~h)
5. ~Kc

Here the first premise is simply Closure 11. Set aside for now how exactly to understand "K*" in this principle and reflect on other features of the argument. The second premise is the skeptics now strengthened central claim – one is not even positioned to know the denial of the skeptical hypothesis. As noted earlier, as one weakens the interpretation of "K*", this key skeptical claim gets stronger, but assume for now that the skeptic will find a good defense of this stronger claim. If the skeptic can find an interpretation of K* on which the closure principle looks strong and P2 is still defensible, the skeptic's overall argument looks promising. Premise 4 in this argument sketch modestly claims that agents are "positioned" to know these entailment facts and does not claim that they are known by all agents. Here the skeptic might promisingly appeal to the fact that the entailments in question are straightforwardly recognizable by anyone with the relevant concepts. The entailments in question are simple conceptual matters: anyone reflecting on the relevant propositions will easily see and come to know that, for example, there being a chair in front of me entails the denial of any skeptical hypothesis explicitly stating that there are no chairs.

Because the sketch of the skeptical argument using Closure 11 is so clean, we think the skeptic's best move is to stick with the form of Closure 11 and attempt to avoid the problems we have pointed out for Closure 11 above by reinterpreting "K*". We have considered and rejected the "all but belief" interpretation of K* which interpreted K* as a conditional ("K*p" read as "Bp $\rightarrow$ Kp" with the "$\rightarrow$" read as either strict, counterfactual, or material implication). What alternatives are available?

Two possible readings appear worth investigating. One might want to consider various diachronic conditional readings of K*. And one might want to consider various "all (for knowledge) but X" readings weaker than the "all but belief" reading considered above. Readings of this latter sort analyze "position to know" as some sort of regular conditional (strict, counterfactual, or material) but this time with a fuller antecedent. Instead of "K*q" being read as "Bq $\rightarrow$ Kq", on this new suggestion K*q is read as "(Bq & X) $\rightarrow$ Kq" with some suitably chosen substitution for X. We now consider both of these possibilities, beginning with the diachronic readings.

The simplest diachronic reading takes "K*q" to be short for "if the agent now adds q to her beliefs, she will thereby come to know q". This simple diachronic reading faces obvious difficulties related to the issues discussed in detail in Section III concerning the basing relation: whether our agent knows q upon acquiring this belief depends upon, among other things, the specific reasons for which belief in q was adopted. Consider an agent who knows both p and that p entails q. If this agent then comes to believe q but on the basis of wholly defective reasons, the agent will not thereby come to know q. This problem should be familiar by now, as should the obvious strategy for attempting to address the problem. The obvious strategy for dealing with this problem is to enrich the diachronic reading in ways analogous to the antecedents of Closure 6 and Closure 7. SK*q, on this enriched reading, would be read as something like: "if S now comes to believe q via and based upon proper inferential grounds from S's beliefs that p and that (p entails q), then S knows q". We have two critical observations about this proposed enriched diachronic reading of "K*".

First, diachronically adding a belief in q to one's belief set, i.e., adding the belief q, might have contingent psychological effects that cause trouble for one's overall epistemic state. Perhaps for at least some agent, the addition of Bq brings with it additional epistemic baggage that eliminates knowledge that p or prevents one from knowing or even being positioned to know that p entails q. Adding Bq even in the right sort of way might, across the time involved in a diachronic reading of this exercise, destroy or defeat one's evidence for q (and/or for "p entails q"). The problem is that closure principles are supposed to be necessary truths, but the claims made by closure principles employing the diachronic interpretation of K* depend on contingent facts about the psychological make-up of agents. It's hard to see how diachronically interpreted closure principles can withstand this criticism.

Second, though this specific diachronic reading of K*q fits well in the consequent of principles like Closure 10 (with K* occurring only once in the consequent) it provides no *general* understanding of "positioned to know". Surely there are ways to be positioned to know things that don't involve deduction from premise beliefs of the form p and (p entails q). This point comes out clearly in considering the enriched diachronic reading of both the antecedent and consequent of Closure 11. The reading doesn't make good sense of the consequent of Closure 11 because it assumes that the agent actually believes (p entails q) while the antecedent of the principle provides only for the agent's being positioned to know this entailment. If our agent does not even believe that p entails q, she can hardly come to know it by deducing it from, among other things, this belief she does not have.

Perhaps this difficulty can be avoided by envisaging a multi-step diachronic process along these lines: if the agent now adds (p entails q) to her belief set and then adds q to her belief set by deducing it from her beliefs that p and that (p entails q), she will then know that q. This suggestion helps with the problem of the previous paragraph but leaves us with a bifurcated understanding of "position to know". After all, this suggestion cannot apply to the reading of K* in the antecedent of Closure 11: apparently that appearance of "K*" must be read in the original way (if she now adds belief in the entailment to her belief set[32] she will then know it). The interpretation of Closure 11, on this bifurcated understanding of K*, will be true if and only if the following multi-step diachronic conditional is true:

Necessarily, IF S is such that she knows p and such that (if she now adds (p entails q) to her belief set she will know that p entails q) THEN S is such that [if she does now add (p entails q) to her belief set and then also adds q to her belief set based on deduction from her belief that p and her belief that (p entails q)] she will then know q.

Though this reading of Closure 11 is far from perspicuous, we do think it is the best available reading for one pursuing what now seems to be an ill-advised diachronic reading of "positioned to know". This multi-step diachronic reading only makes the first problem we discussed for the enriched diachronic approach more serious. On this reading, there are multiple steps at which various things can go epistemically wrong in one's attempt to reach knowledge that q (defeaters can be acquired, evidence lost, etc…). It is quite implausible to suggest that no actual or possible agent's path to Kq would be blocked somewhere along the way in one of these ways.[33] We see no hope for finding a plausible weakened consequent closure principle involving diachronic conditionals.

We turn now to regular conditional readings of "position to know" but this time with an eye on readings weaker than the "all but belief" reading discussed above. We now consider reading "position to know" as some type of regular conditional with a fuller antecedent than simply "Bq" – we examine readings that read "K*q" as "(Bq & X) $\rightarrow$ Kq" where X is some specified additional conditions on believing q. We will consider three substitutions for X, all somewhat familiar from our discussion in Section III:

(i) "K*q" means "Bq based on deduction from p and (p entails q) $\rightarrow$ Kq"
(ii) "K*q" means "Bq solely based on deduction from p and (p entails q) $\rightarrow$ Kq"
(iii) "K*q" means "Bq solely based on deduction from p and (p entails q) and 'nothing else wrong' $\rightarrow$ Kq"

Note first that just as in the discussion of enriched diachronic conditional readings of K*, none of these suggested readings of "K*" permits a univocal understanding of "K*" as it appears in both the consequent **and** antecedent of Closure 11. We have to leave the interpretation of K* in the antecedent of the principle largely unanalyzed.[34] This amounts to presenting the following familiar skeptical argument sketch, with intuitive defenses of the premises inserted:

P1.	(K*p* and K*(*p* entails *q*)) $\rightarrow$ K**q*	Closure 11
P2.	~K*~h	Can't Know This
3.	~Kc or ~K*(c entails ~h)	
P4.	K*(c entails ~h)	Easy to know this (unanalyzed K*)

5. ~Kc

P1 here is simply Closure 11. Premise 2 is the standard skeptical pivot point. P4 is not interpreted using any of the "all but belief and X" understandings of K* now in play. Instead the "position to know" locution at this step in the argument is left unanalyzed and the premise is defended with an intuitive appeal to how easy it is to know or at least be positioned to know the straightforward entailments relation specified in P4.[35]

Note second that the material conditional readings of (i), (ii), and (iii) are strictly equivalent to Closures 6, 7, and 8 from Section III. We here exploit the simple logical transformation used earlier to show that material conditional readings of the "all but belief" understanding of K* are strictly equivalent to already discussed strengthened antecedent closure principles. Because, as shown in Section III, none of Closures 6-8 are of use in skeptical argumentation, it immediately follows that the material conditional interpretations of (i – iii) are similarly useless: if one principle is useless in a skeptical argument then so is any principle strictly equivalent to it.

The more serious readings (i), (ii), and (iii) are the modal conditional readings. For what should be familiar by now reasons, the counterfactual and strict conditional readings of (i) and (ii), however, respectively imply the truth of Closure 6 and Closure 7.[36] We claimed in section III that Closures 6 and 7 and are not merely demonstrably useless to the skeptic but also false. We focused our argument in that section on the former point and so will now remind the reader of the main points we sketched to make the claim that Closures 6 and 7 are not only useless to skeptic but also false.[37]

Closure 6 has trouble with overdetermination cases. Where S believes q both for intuitively good reasons (e.g., because of deduction from known "premise beliefs" of the form p and (p entails q)) and bad reasons (e.g, because it's Tuesday or because a wholly unreliable witness testified to the truth of q) but believes q most firmly and confidently on the basis of the bad reasons it seems implausible to suggest that S knows that q. Closure 7 seems to have serious trouble at least with cases involving "global defeaters" – the most relevant case is one in which agent S satisfies the antecedent of Closure 7 but additionally believes (on the basis of apparently powerful evidence) that deduction is a horrible way, from the epistemic point of view, to add beliefs. Perhaps S believes that this is true of deduction at least on this occasion (the occasion of adding q to his belief set). Perhaps S's normally quite reliable Doctor has informed S that his deductive reasoning powers should not be trusted for the next hour because of medication S has taken and the Doctor supports this by showing S some simple logic problems the Doctor administered to S the last time S took the medication in question.

If Closures 6 and 7 are not only useless to the skeptic but also simply *false*, as we claim, then of course any Closure principle implying either Closure 6 or Closure 7 is also false and is therefore of no use in a skeptical argument. The subjunctive and strict conditional readings of (I) and (II) above imply, respectively, Closure 6 and Closure 7. These principles are therefore false and of no use in a skeptical argument.

The modal conditional readings of (III), however, cannot be criticized in this way. While the material conditional reading of (III) is equivalent to Closure 8 and therefore of no use in a skeptical argument, the modal conditional readings of (III) imply, but are not implied by, Closure 8. Though closure principles equivalent to useless principles like Closure 8 must themselves be useless, it is not the case that principles implying useless principles must be useless. Additionally, as we pointed out in Section III, Closure 8, though imprecise, must be true: we therefore can't argue for the uselessness of modal conditional readings of (III) on grounds that they imply a false closure principle. We therefore need an alternative criticism of the modal conditional readings of (III).

Fortunately, a serious criticism of these readings of (III) is immediately available. This interpretation of "position to know" makes it *far too easy* to be "positioned to know something". In fact, these readings of (III) as interpretations of "position to know" make it maximally easy to be positioned to know things. Consider the strict conditional reading of (III) -- this reading implies says that I am positioned to know that q if it's true that:

Necessarily, if I to believe q solely on the basis of deduction from p and (p entails q) and nothing else go wrong then I know q

(the counterfactual reading replaces the necessity operator with counterfactual "were I to believe q…" in the antecedent and "I would know q" in the consequent). Because "nothing else go wrong" is being read in such a way as to bracket all plausible counterexamples to principles like Closure 7, it seems that on this reading one is going to be ruled to be "positioned to know" any proposition at all. Most notably, one is going to be judged to be "positioned to know" some q even when it's true that were one to believe q it would be on the basis of bad reasons, in the face of multiple sources of defeat, and despite a global defeating belief concerning the method of belief formation in question.

The powerful antecedent of this reading ("nothing goes wrong") abstracts away from the sometimes messy real epistemic contexts that agents occasionally find themselves in. Sometimes we fail to be "positioned to know" things precisely because our overall epistemic context is problematic (defeaters are present, our belief forming mechanisms aren't working well, etc…). The modal conditional readings of (III) abstract away from these problems and rule us positioned to know propositions we clearly are not positioned to know. These readings are therefore inadequate readings and cannot be of use in a skeptical argument.

In addition to the simple inadequacy of the proposed understanding of "positioned to know", there is a second serious problem for this reading when considering the skeptical argument. On this weak understanding of "positioned to know", the skeptics claim (at P2 in the standard argument sketch) that one is not even "positioned to know" the denial of skeptical hypothesis becomes highly suspect and maximally deniable. Anti-skeptics inclined to resist the standard skeptical claim that we don't know the denial of the skeptical hypothesis will, as noted above, be even tempted to deny the claim that we are not "positioned to know" this denial and they will justifiably be maximally tempted to deny this claim on this incredibly weak understanding of what it is to be positioned to know something.

We see no hope for the skeptic in these various fuller antecedent interpretations of "K*" and having reached this pessimistic conclusion we have also exhausted the full range of natural interpretations of "position to know" that we have identified. Perhaps this represents a failure of imagination on our part, however, and not a failure for this skeptical strategy. We close this section with a discussion of a final skeptical possibility involving our favored weakened consequent principle, Closure 11.

Skeptics will surely have seen the difficulty with the strengthened antecedent principles of Section III. Because of this, they will rightly focus on the weakened consequent principles of the present section. Though we have not been able to identify a plausible and skeptic friendly interpretation of K*, we think that skeptics might want to press on with the sketch of a skeptical argument involving Closure 11 presented above. Here's that argument sketch again:

P1.	(K*p* and K*(*p* entails *q*)) → K**q*	Closure 11
P2.	~K*~h	Can't Know This
3.	~Kc or ~K*(c entails ~h)	
P4.	K*(c entails ~h)	Easy to know this (unanalyzed K*)
5.	~Kc	

We noted above that the skeptic will almost certainly have to make use of different understandings of K* in the antecedent and consequent of the first premise of this argument (and hence in P2 and P4 of the argument as well). We noted further that the natural substantive suggestions for interpretations of K* in this argument are really suggestions for interpreting the K* of the consequent. The K* of the antecedent of P1 (and P4) was to be left unanalyzed. This did not seem to present an immediately devastating problem for the skeptic because intuitively the conceptual matters of concern in the antecedent of the principle are propositions that all agents with the relevant concepts are at least plausibly thought to be "positioned to know".[38]

In light of the troubles we have revealed in attempting to find an analysis of the K* of the consequent of the principle (Closure 11), the skeptic should more fully embrace this "no analysis" approach to the principle. Maybe the skeptic should attempt to run the skeptical argument without an analysis of any of the K*s of the argument. There is, we grant some intuitive appeal to this approach. This intuitive appeal is revealed in the support sketched above for P2 and P4 of the argument.

The intuitive and not wholly implausible idea behind P2 is clearly that the skeptical hypothesis is well chosen so that one cannot know its denial (or that it's at least very hard to know its denial). We aren't positioned to know this denial because it's so hard to know it. By contrast, the intuitive and not wholly implausible support for P4 is that these straightforward conceptual matters (about the relation between, for example, the chair hypothesis and the "no chair" skeptical hypothesis) are maximally easy to know. We are positioned to know these entailment relations because they are so easy to know.

With that intuitive "no analysis" support of the premises in place, the skeptic might seem to be well-positioned in defending the general skeptical argument. This verdict is premature however. We still must evaluate the principle driving the argument (P1 in the argument sketch). This is not easy to do in the absence of a specification of the interpretation of the two "K*s" in the principle but we think the skeptic's general idea here is clear and, once again, not wholly implausible. The skeptic's idea must be the following. If one knows that p and is rather well positioned to know that p entails q (because it's a simple entailment relation) then of course one is positioned to know q: after all, one well positioned to know the entailment of q by p and already knowing that p could simply put these two "premise beliefs" together and infer (and thereby come to know) that q.

We see no other way for the skeptic to offer purely intuitive support for P1 of the skeptical argument sketch (the "no analysis" version of Closure 11). Recall that in the absence of a specification of the meaning of "positioned to know" this is the only sort of support to which the skeptic can appeal. This support, intuitive though it might be, is inadequate to the skeptic's task. The skeptic appeals to the intuition that one knowing p and knowing (or easily being able to know) that p entails q could simply put these bits of knowledge together and come to know q. On this basis the skeptic endorses the intuitive reading of Closure 11. Note, however, that this support is tantamount to endorsing principles like Closure 6 and/or Closure 7 (the imprecision about which principle is involved is simply a function of the informality of the skeptical reasoning being discussed). Closure 6 and Closure 7 are false and so are not available in an intuitive defense of the “no analysis” version of Closure 11. We therefore reject this “no analysis” defense of the “positioned to know” version of the skeptical argument.

Having rejected this final skeptical strategy, we see no further way for the skeptic to offer a plausible skeptical argument using weakened consequent knowledge-closure principle. We see nothing for non-skeptics to fear in this exploration of closure-based skeptical arguments that attempt to find plausible closure principles by weakening the consequent of implausible simple knowledge closure principles.

V. Conclusion

The best sketches of closure-based skeptical arguments appear to present non-skeptics with the beginnings of a serious challenge to typical knowledge claims. It is uncontroversially true, however, that the knowledge-closure principles employed in these simple argument sketches are not sufficient for use in a general and powerful skeptical argument. We have explored a wide range of possible and natural attempts to produce plausible knowledge-closure principles that might be of service in such a skeptical argument. It is not easy to find plausible principles of this type (let alone plausible principles of use to the skeptic). Though we did manage to find a few plausible principles, these principles are clearly not of use in skeptical argumentation.

Principles discussed in Section III (the strengthened antecedent principles) are the easiest to generate, understand, modify and defend. But these principles appear quite clearly to be of no use in a general and powerful skeptical argument. Principles discussed in Section IV (the weakened consequent principles) are not as easy to understand, generate and defend. This is because of ambiguities and uncertainty about plausible analyses of the key notion of being “positioned to know” something. Principles of this sort are, from a structural point of view, more promising principles for use in general skeptical argumentation. But upon closer analysis these principles tend to be much more closely connected to their strengthened antecedent cousins than an initial inspection would reveal. The connections we have explored between the two kinds of modifications of simple closure principles lead us to the conclusion that weakened consequent knowledge-closure principles are also of no sue to the skeptic.

We conclude that non-skeptics have nothing at all to fear from knowledge-closure based skeptical arguments.[39]

References

Brueckner, A. L. (1985): “Skepticism and Epistemic Closure”, *Philsophical Topics* 13, 89-117.

Brueckner, A. L. (1994): “The Structure of the Skeptical Argument”, *Philosophy and Phenomenological Research* 54, 827-35.

Cohen, S. (1998): “Two Kinds of Skeptical Argument”, *Philosophy and Phenomenological Research* 58, 143-59.

Cohen, S. (1999): “Contextualism, Skepticism, and the Structure of Reasons”, *Philosophical Perspectives 13: Epistemology*, 57-89.

David, M. and T. Warfield (forthcoming) …. “Six Possible…”

DeRose, K. (1995): “Solving the Skeptical Problem”, *Philosophical Review* 104, 1-52.

Devitt, M. (1996): *Coming to our Senses*, Cambridge: Cambridge Univ. Press.

Dretske, F. (1970): “Epistemic Operators”, *The Journal of Philosophy* 67, 1007-23.

Feldman, F. (1995): “In Defence of Closure”, *The Philosophical Quarterly* 10, 487-94.

Harman, G. (1973): *Thought*, Princeton, Princeton Univ. Press.

Klein, P. (1981): *Certainty: A Refutation of Skepticism*, Minneapolis: University of Minnesota Press.

Klein, P. (1992): “Scepticism, Contemporary”, in J. Dancy and E. Sosa, eds., *A Companion to Epistemology*, Oxford: Blackwell, 458-62.

McGee, V. (1985): “A Counterexample to Modus Ponens”, *The Journal of Philosophy* 82, 462-71.

MAITZEN

Moore, G. E. (1939): “Proof of an External World”; reprinted in *Philosophical Papers*, London: Allen and Unwin 1959.

Nozick, R. (1981): *Philosophical Explanations*, Cambridge: Harvard Univ. Press.

VanArragon, R. (2001): *Externalism, Skepticism, and Relevant Alternatives*, University of Notre Dame, Doctoral Dissertation.

Vogel, J. (1990): “Are there Counterexamples to the Closure Principle?”, in M.D. Roth and G. Ross, eds., *Doubting*, Dordrecht: Kluwer Academic Publishers, 13-27.

VOGEL, J.

Warfield T. (2004): “When Closure Does and Does Not Fail: A Lesson From the History of Epistemology,” *Analysis*.

MORE REFERENCES TO BE ADDED

Notes

[1] One apparently alternative defense of the second premise invokes an *undertermination* principle, making use of the idea that Leopold's evidence underdetermines the "chair hypothesis" vis-à-vis the skeptical hypothesis; see Brueckner 1994, and Cohen 1998. VOGEL?

[2] Nozick himself argued, implausibly, that closure principles fail as a consequence of details of his analysis of knowledge; see his 1981, chap. 3.2. This argument is no more plausible than Nozick's widely rejected analysis of knowledge. Additionally, the details of Nozick's argument that Closure fails given his account of knowledge are mistaken. See Warfield 2004 for a discussion of this issue.

[3] See Klein 1981; Brueckner 1985.

[4] Cf. Nozick 1981, pp. 47f.; Vogel 1990, p. 13; Cohen 1999, p. 62.

[5] This assumes, as is plausible, that knowledge of a conjunction entails knowledge of each conjunct. Hence, the *antecedent* of Closure 3 entails the *antecedent* of Closure 2; hence, Closure 2 entails Closure 3. For any two principles with the same consequent, if the antecedent of one entails the antecedent of the other, then the other entails the one.

[6] REFERENCES!

[7] Because the antecedent of Closure 2 trivially entails the antecedent of Closure 4; see note 5.

[8] Though if one preferred a variant of Closure 6 that employs a single K-operator in the manner of Closure 3, then reading (ii) would seem the better choice.

[9] The conditional in this example is not obviously a strict conditional. We occasionally use examples of this form where in our judgment nothing of consequence turns on the precise modal strength of the conditional. Examples more clearly involving strict conditionals can easily be substituted at a cost of additional complexity in the presentation.

[10] See VanArragon 2001, chap. 2.

[11] They might be students of Vann McGee, the author of "A Counterexample to Modus Ponens"; see McGee 1985.

[12] See MAITZEN for an extensive discussion of closure and epistemic paradox. In David and Warfield (forthcoming) we discuss these and some other challenges; we will return to them a bit in Section IV.

[13] This is somewhat disingenuous. After all, the philosophical skeptic is (usually) not a real person but rather the metaphorical personification of our own skeptical worries. Consequently, the premises of a serious skeptical argument must be premises that we ourselves are initially inclined to accept: in the end, the skeptic's tasks are our tasks.

[14] The skeptical hypothesis Descartes considers in the first Meditation appears to be a general hypothesis much like H; cf. *Meditations*, AT VII, p. 21.

[15] What about other candidates for playing the role of the general skeptical hypothesis? They are not easy to find. Take the hypothesis that we are all dreaming: it's not easy to come up with any perceptual beliefs that *entail* the denial of this hypothesis. Or take the hypothesis that there is an evil demon who makes it the case that all our perceptual beliefs are false—call it H1: "There is an evil demon who makes it the case that ($\forall S, \forall p \in$C: if SBp, then p is false)." Our perceptual beliefs, beliefs in propositions of class C, won't entail the denial of H1. To see this, consider a particular proposition, c, say the proposition that there are chairs in the room, and assume that it belongs to C and that some person S believes it. Proposition c of course entails that c isn't false. But that's not enough. To entail the denial of H1, c would have to entail the denial of the conditional "if c∈C and SBc, then c is false"; that is, it would have to entail the conjunction "c∈C and SBc and c is not false"; but it doesn't. Take, finally, the hypothesis that we are brains in a vat. This skeptical hypothesis is significantly less general than H: only perceptual beliefs entailing that our bodies are more than just our brains will entail the denial of the brains-in-a-vat hypothesis. Anyway, the problems we are about to raise for skeptical arguments employing H also arise for arguments employing this considerably less general skeptical hypothesis.

[16] For some authors whose remarks indicate that they tend to think of closure based skepticism in terms of the strategy of particular skeptical hypotheses see: Dretske 1970; Nozick 1981, pp. 204-11; Klein 1992, p. 459; Feldman 1995, p. 487; Cohen 1999, p. 85, fn. 31.

[17] The already somewhat tortuous yet still rather imprecise characterization of C given earlier (see the paragraphs following the Argument Sketch) is not, it seems, good enough to forestall this objection.

[18] We apologize for this awkward rendering of CA. It is tempting to paraphrase CA as saying, roughly, that agents who satisfy the antecedent have *no* perceptual knowledge. But note, there is no quantifier corresponding to the "no" in CA's consequent. CA's consequent applies to agents and propositions that pairwise satisfy the antecedent: it's not

easy to put this fluently. Also, since the full CA is a universal generalization, it's not quite right to talk about its antecedent/consequent. What this talk refers to is the antecedent/consequent of the part of CA that is explicitly exhibited above.

[19] Getting the relevant "further premise" would require the skeptic to show for all agents *S* and all propositions *p*, such that that $p \in$C and *S*B*p*: *if* it is *not* the case that *S*K(*p* entails ~H) and *S*B~H bod from *p* and (*p* entails ~H), *then* ~*S*K*p*.

[20] For an author who thinks of closure-based skepticism along the lines of Argument C, see Brueckner 1985, pp. 89f.

[21] But even such decompositional inferences must be handled with care: the proposition that John spends counterfeit money does not analytically entail the proposition that John spends money.

[22] It is noteworthy that Moore in his "Proof of an External World" (1939), which is well-known primarily for its last few pages, actually spends most of his time on what appears to be an attempt to show that the proposition that this is a hand analytically entails the proposition that there is an external world. Moore, it seems, assumed that the skeptic would pursue the two-part strategy of first employing Argument C and then Argument A but with Closure 1A instead of Closure 6.

[23] For example, Devitt (1996) argues, roughly, that analytic entailment should rather be understood in terms of a disposition to infer the entailed proposition. On this construal, Closure 1A will not be of much help to the skeptic presently under consideration.

[24] Equivalently: necessarily, for all agents S, and all propositions *p* and *q*, if S knows *p* and if *p* entails *q*, then S is positioned to know *q*.

[25] We set this promising anti-skeptical move aside not because we are not willing to make it – we do think that a promising anti-skeptical strategy could be found by arguing that we know (or are at least positioned to know) the denial of the skeptical hypothesis. We set this issue aside in order to keep our focus on those parts of the skeptical argument intimately featuring knowledge-closure principles.

[26] This second point might not do justice to all possible understandings of the "positioned to know" locution. Perhaps just as I am poised to enter a room I am "positioned to know" that someone is in the room even though no one is currently in the room (my being in the room upon entry will imply that someone is in the room and upon entry I'll be ideally positioned to know that I'm in the room). As a part of this discussion of knowledge-closure principles, however, we don't see any reason to worry about examples like this one. If for some reason we do not see we should worry about examples like this one, we could modify our claim to read: one is positioned to know only what is (or would be if/when believed) true.

[27] Using "warrant" in Plantinga's sense, one could say that S is positioned to know *p* iff *p* is for S a warranted truth (whether or not SBp). We take no stand here on whether "warrant" implies "truth" – if it does the explicit inclusion of truth in this claim does no harm. Strictly speaking (because knowing is a way of being positioned to know on our understanding) K*p = Kp or K minus B.

[28] Cf. Cohen 1999, p.84 note 14; Klein 2000, p.110.

[29] For this suggestion to be plausible, the conditional here must presumably be read with *at least* the modal force of a counterfactual: no one would suggest that you are positioned to know every truth that you fail to believe. In the text we consider reading the conditional as a strict conditional which seems to be the most natural reading. The formal points made in the text stand whether the conditional is interpreted as a strict conditional or a counterfactual.

[30] Whether the principles are equivalent to or imply the previously encountered principles depends on the strength of the modality in the K* conditional. Reading them in any plausible way (as a strict or counterfactual conditional) gives us the same results: Closure 9 implies the problematic Closure 9'' and Closure 10 implies the similarly problematic Closure 5. Reading the conditionals in question (implausibly) as material conditionals would make our criticism even stronger: on this reading the paired weakened consequent and strengthened antecedent conditionals are strictly equivalent.

[31] The "at least" parentheticals are redundant reminders that we are understanding "positioned to know" in a way implying that one way of being positioned to know something is to know it.

[32] The agent would presumably be required to do this on the basis of non-defective conceptual reasoning. We set aside this difficulty though it is not a trivial complication for the skeptic.

[33] Recall that in the full dress version of the skeptical argument, the argument will feature a necessity operator out front and universal quantification over all agents.

[34] We note only that it will have to be understood as something like "either the agent knows the entailment or would believe it solely on the basis of good and sufficient reasons (and therefore would know it)". We do not think the skeptic is on terrible ground here given that knowledge of the relevant entailments should be rather easily and cleanly a priori accessible even to agents who haven't previously thought about the skeptical hypotheses.

[35] The point would be generalized in defense of the relevant full dress universally quantified premise of the more formal skeptical argument of which this is a sketch.

[36] They imply the material conditional readings of (i) and (ii) which are equivalent to Closure 6 and 7 respectively.

[37] See David and Warfield forthcoming for fuller arguments for the falsity of Closure 6 and Closure 7.

[38] Might be too permissive – NOTE agent confused temporarily about conceptual matters…

[39] Early versions of various parts of this paper were presented in colloquia at the University of Rochester, University of Cincinnati, Florida State University, University of Notre Dame, and University of Missouri. We thank our audiences and hosts on these occasions. We have received generous feedback on this project from many philosophers and apologize for not thanking all of them by name here (the list would be longer than the paper!). We do want to especially thank EJ Coffman, Earl Conee, Thomas Crisp, Richard Feldman, and Leopold Stubenberg for helpful feedback and valuable advice.