Epistemic Closure Principles

Steven D. Hales
Bloomsburg University

Much has been made of the principle that knowledge is closed under known implication, or put another way, that knowledge is transmitted or flows down through known implication. And for good reason—the principle is a very plausible one. It commands such respect that Richard Feldman writes, "some version of the closure principle ... is surely true. Indeed, the idea that this principle is false strikes me, and many other philosophers, as one of the least plausible ideas to come down the philosophical pike in recent years" (Feldman 1994, 1). If the principle *is* true, then it constitutes a valuable discovery for epistemology. Principles like it have been used by foundationalists to show how derivative knowledge is based on the immediately known foundations, Gettier relied on something like it to undermine knowledge as justified true belief, and skeptics are said to depend on it to make their case against empirical knowledge. Moreover, such a principle would allow us to deduce what someone knows from other things that they know. That is to say, deduce in the strong sense—closure does not merely provide for inferring probable beliefs or knowledge, but for deducing the existence or presence of these with all the force of logic itself. This is one of the great promissory notes that doxastic logic has issued to epistemology.

Unfortunately, this note will never be made good. As I will argue, the closure principle is false. What's more, several related closure principles are false as well; or at least they seem to be false with respect to actual knowers. It is not too hard to make them come out true if we only apply them to perfectly rational knowers or other specially tailored fictional creatures.[1] But their inapplicability to real people makes it difficult to see what benefit doxastic logic might have for an epistemology devoted to actual knowers in the actual world.

Steven D. Hales received his Ph.D. from Brown University in 1992. He is Assistant Professor of Philosophy at Bloomsburg University. Recent articles have appeared in Synthese, Canadian Journal of Philosophy, *and* History of Philosophy Quarterly.

The only closure principle that seems to turn out true for actual knowers may be formally trivial, and is certainly worthless for the project of deriving what someone believes or knows from other things they believe or know. Thus, in this essay I will argue that several closure principles are false, and that their falsity shows doxastic logic to be as otiose for epistemology as Carnap's "perfect" logical language is for Gamut's semantics of natural language.[2]

Closure is a property of sets. A set S is closed under a relation R just in case every element of the set is such that anything it is R-related to is a member of S. In the case of knowledge, the set is that of items of knowledge. One sort of closure principle is that of closure under implication, and can be expressed as follows.

K1.　❑ $\forall$ x $\forall$ y ((Kx & (x $\rightarrow$ y)) $\rightarrow$ Ky)
　　　Closure of knowledge under implication[3]

That is, for any two propositions x and y, if you know that x, and x implies y, then you know that y. This principle is most often associated with the work of Jaakko Hintikka (cf. Hintikka 1962). Many philosophers have noted that this principle is false—it is just too strong. K1 is a sort of omniscience principle; it commits knowers to knowing all of the logical consequences of their knowledge. It is easy to see, especially in the case of mathematical knowledge, why this does not hold. There are a variety of ways to modify the principle so as to make it more plausible. Here are a couple of modified versions:

K2.　❑ $\forall$ x $\forall$ y ((Kx & B(x $\rightarrow$ y)) $\rightarrow$ Ky)
　　　Closure of knowledge under believed implication

K3:　❑ $\forall$ x $\forall$ y ((Kx & J(x $\rightarrow$ y)) $\rightarrow$ Ky)
　　　Closure of knowledge under justified implication

Both of these principles are also false for obvious reasons. One reason is that truth is not guaranteed to the content of the consequent. For example, I might know that the sun is shining, believe that this implies that the birds are singing, and still fail to know that the birds are singing. After all, they might be silent despite my beliefs. Perhaps the most famous of the closure principles, and the central focus of this essay, is the following:

K4:　❑ $\forall$ x $\forall$ y ((Kx & K(x $\rightarrow$ y)) $\rightarrow$ Ky)
　　　Closure of knowledge under known implication

That is, for any two propositions x and y, if you know that x, and you know that x implies y, then you also know that y. K4

has a clear advantage over K2 and K3—truth is guaranteed to the content of the consequent. Among those who accept K4 are Jonathan Vogel (Vogel 1990),[4] Gail Stine (Stine 1976), and Anthony Brueckner (Brueckner 1985a, 1985b). Others, to be discussed later, are sympathetic to K4 and explicitly advocate similar principles. Foes of the principle include Fred Dretske (Dretske 1970), Robert Nozick (Nozick 1981), and Colin McGinn (McGinn 1984). The criticisms of Dretske and Nozick rest upon their specific analyses of knowledge. For example, Nozick's account of knowledge is roughly this: S knows that p iff p is true, S believes that p, if p were not true then S would not believe p, and if p were true then S would believe p (Nozick 1981, 179). While I will not rehearse his proof here, it just turns out that knowledge is not closed under known implication if we define knowledge this way. Clearly what is wanted is a pretheoretical argument against K4. A fervent defender of K4 might look at Nozick's results against K4 and decide so much the worse for Nozick's analysis of knowledge! But the two counterexamples I will give (and it is reasonably clear how to construct more in the same vein) rely only on our pretheoretic intuitions about knowledge. Of the three foes of K4, only McGinn claims to provide a pretheoretical critique of the principle. Let us begin, then, with an examination of his argument. He writes,

> The following seems an intuitively correct principle: one can *know* that p only if one can *tell whether p*—I can know that (e.g.) it is raining outside only if I can tell whether it is raining outside. Let us apply this principle to my putative knowledge that there is a table in front of me and that I am not a brain in a vat. Can I tell whether there is a table there? I think that in the ordinary use of the phrase 'tell whether', what this requires is that I can distinguish there being a table from there being a chair or a dog or some such. So, granted that conditions are normal—there is a table there, my eyes are functioning normally, etc.—I *can* tell whether there is a table there. But can I tell whether I am a brain in a vat? ... [W]hat is required for telling whether I am a brain in a vat is that I be able to *distinguish* my being a brain in a vat from my not being a brain in a vat. But it seems clear that I lack this ability—I cannot tell whether I am a brain in a vat because I have no means of distinguishing being in that condition from not being in that condition. (McGinn 1984, 543)

Thus, I can consistently know that there is a table in front of me, know that this entails that I am not a brain in a vat, and not know that I am not a brain in a vat. Hence nonclosure. It is clear that McGinn's argument rests on his "intuitively correct" principle, and so if the principle is false, or at least unintuitive, then his argument against K4 will prove unsatis-

factory. McGinn's position is that a necessary condition for my knowing *p* is that I am able to *tell whether p* (is true). Preanalytically there seems to be a variety of methods that one could employ to tell whether something is true. One perfectly decent way is deduction itself. Surely if I deduce *p* from other propositions, this would count as a way of telling whether *p*. Assume McGinn is right that I cannot perceptually distinguish being an envatted brain from not being one—nevertheless such distinguishing is only one form of *telling whether*. Suppose I claim to know that I am not an envatted brain by deducing this from other bits of knowledge that I have. For example, I know that I am in Pennsylvania, know that this entails I am not a brain in a vat near Alpha Centauri, and thereby deduce that I am not such a brain in a vat. This ought to count as telling whether I am a brain in a vat. McGinn might deny that this is a legitimate procedure, since it seems to presuppose the truth of K4, the very principle at issue. However, I maintain that McGinn would beg the question with such a response. We are currently considering deduction to count as *telling whether*, and this argument seems preanalytically like a legitimate use of deduction. I conclude that McGinn has not produced a satisfactory preanalytic argument against K4.[5]

Another unpromising way of arguing against K4 is suggested by Anthony Brueckner's claim that "knowledge is closed under known implication only if each necessary condition for knowing is so closed" (Brueckner 1985a, 91; the claim is repeated in Brueckner 1994, 831). Thus, if we can merely show that, say, belief is not closed under known implication, we can conclude that knowledge is not either. However, Brueckner's claim, for which he provides no argument, commits the informal fallacy of division. The set of items of knowledge is a proper subset of the set of beliefs. Suppose that the set of beliefs is not closed under known implication. This does not entail that every proper subset of the set of beliefs is not closed. Perhaps what you add to belief to get knowledge is precisely what is needed to eliminate those cases that exhibit nonclosure. Put metaphorically, knowledge is a finer propositional sieve than belief, and winnows out things that mere belief does not. Thus, if we are to argue against K4, we must aim directly at knowledge itself, not at its components. It is obvious what form a counterexample must take—K4 is offered as a necessarily true conditional. It is defeated if it is possible for its antecedent to be true and yet its consequent be false.

Here is the first counterexample:

Suppose a logic student is told in class that "it is strange but true, logically true, for any two people John and Mary, that there is someone who, if admired by John, admires Mary."[6] This student refuses to accept this bizarre claim, and

is presented by the professor with the following proof: Consider this sentence: (i) $\exists z(Fxz \rightarrow Fzy)$. There are two ways of substituting the free variables for the bound existential variable in (i). They yield 'Fxx→Fxy' and 'Fxy→Fyy'. A little reflection reveals that the disjunction 'Fxx→Fxy ∨ Fxy→Fyy' is valid—i.e., true under all interpretations. Each disjunct implies (i); thus, any interpretation that makes either half of the disjunction true makes sentence (i) true. The disjunction implies (i), and since the disjunction is valid, so is (i). Interpreting 'Fxy' as 'x admires y', and allowing the constants 'x' and 'y' to stand for John and Mary respectively, we can see that for any two people John and Mary, there is someone who, if admired by John, admires Mary.

Consider the student after having been presented with this proof. Let us assume that the student believes the premises to be true, and moreover, knows that they are true. If we are prepared to grant that logic students do sometimes learn material presented in class, then this is surely reasonable. Aside from knowing the truth of the premises, suppose that the student also knows that they entail the conclusion. Now, given the truth of the closure principle K4, we can conclude from these facts that the student does in fact know it to be a logical truth that for any two people John and Mary, there is someone who, if admired by John, admires Mary. The problem is that it seems completely possible that the student know the premises, know what they entail, and continue to refuse to accept the repugnant conclusion. When he reflects on the conclusion his superstitious and skeptical side takes over. The student need not especially disbelieve the conclusion; he just does not believe it.

It might be objected by a defender of K4 that a refusal to accept the conclusion shows that the student either does not know one of the premises, or does not know that the inference is valid. However, this response either begs the question or reveals brute opposing intuitions about knowledge. Surely we could have plenty of good reasons for ascribing both knowledge of the premises and knowledge that they imply the repugnant conclusion to the student. If we ask him whether his refusal to accept the conclusion means that he is rejecting one of the premises of the argument presented in class, his answer is no. Likewise when asked, the student denies that he means to reject the rules of inference. Not only is the belief condition satisfied, but the other conditions for knowledge (whatever they are) appear to be met as well. That is, the premises are true, they really do imply the conclusion, and the student is properly connected with these truths. He has accepted that they are true on the basis of legitimate authority (his logic professor), or perhaps can reproduce the reasons for thinking that they are true. If the only reason that can be provided for believing that the student does not know the truth of the pre-

mises or that they imply the repugnant conclusion is that to say otherwise violates the closure principle K4, this is only to assume that principle really is true. Since this is precisely the issue at hand, this objection begs the question.

Let us consider another case that has the same conclusion. This is the case of the Monty Hall Puzzler. The Puzzler is as follows. Suppose you are a contestant on the game show *Let's Make a Deal*. The host of the show, Monty Hall, shows you three closed doors, A, B, and C. He tells you that behind one of the doors is a new car (which you want), and behind each of the other two doors is a goat (which you do not want). Monty knows the distribution of prizes behind the doors. You are allowed to choose one of the doors (and the prize behind), with an option to switch to another door and prize. Suppose you choose door A, which remains closed. Monty then intentionally shows you a goat by opening door B. You now have the option of sticking with the door you originally chose, A, or switching your fortunes to door C. If you stick, you get whatever is behind A, and if you switch, you get whatever is behind C. Since you saw a goat behind door B, you know that the car is behind one of the two remaining doors, and a goat is behind the other. The question: should you switch, stick, or is it a matter of indifference?

There is a strong intuition that it is a matter of indifference. After all, there are two remaining doors: one with the car, and one with a goat. Therefore, it seems to be a 50-50 chance either way. However, this initial intuition is misleading. As it turns out, the switch strategy will win the car 2/3 of the time, and the stick strategy will win only 1/3 of the time. Moreover, this is provable from Bayes's Theorem. Here, in fact, is the proof.[7]

> Let h be the hypothesis.
> Let e be the evidence.
> Let $c(x)$ be the chance of x .
> Let $c(x/y)$ be the chance that x is true given that y is.

"$c(h)$" means the chance assigned to the hypothesis before we have any evidence for it—this is the "antecedent" or "prior" chance. "$c(h/e)$" means the chance of the hypothesis now that we have the evidence e. Given this vocabulary, we add Bayes's Theorem:

Bayes's Theorem: $c(h/e) = c(e/h) * c(h)/c(e)$

That is, the chance of the hypothesis given the evidence equals the chance of the evidence given the hypothesis times the quotient of the antecedent chance of the hypothesis divided by the antecedent chance of the evidence. What we want to find out

in the Monty Hall Puzzler is the chance of the hypothesis given the evidence—the hypothesis that the car is behind our original choice of door A, given the evidence that Monty shows us a goat behind door B. So we want to solve for the left side of the equation. We can apply Bayes's Theorem to the Puzzler as follows:

1. c (the car is behind door A/ Monty opens door B) = c (Monty opens door B/ the car is behind door A*)
   c (the car is behind door A) ÷
   c (Monty opens door B)

2. The chance that Monty opens door B given that the car is behind door A is 1/2. This is so because the car's being behind A guarantees that there are goats behind doors B and C. Monty could open either one of these doors to show us a goat, and so there is only a 50% chance for him to open B.

3. The antecedent chance that the car is behind door A is 1/3—it could be behind any of the three doors.

4. The antecedent chance that the Monty opens door B is 1/2; here's why. There are three possible cases. Case one: we choose A and the car is behind A. In this case goats are behind B and C, and Monty could open either of these doors to show us a goat. Thus the chance that he will open B is 1/2. Since the antecedent chance of the car's being behind A is 1/3, the joint chance of the car being behind A and Monty opening B is 1/6. Case two: we choose A and the car is behind B. In this case the chance that Monty opens B is zero, since he will only show us a goat. The joint chance of the car being behind B and Monty opening B is also zero. Case three: we choose A and the car is behind door C. Monty cannot open C (since the car is behind it), and cannot open A (since we have picked it), and can only open B to show us a goat. The chance that he will open B is 1. Since the antecedent chance of the car's being behind C is 1/3, the joint chance of the car being behind C and Monty opening B is 1/3. The overall antecedent chance of Monty opening door B is the sum of the chances of the three possible cases. 1/6 + 0 + 1/3 = 1/2. Hence the antecedent chance of Monty opening B is 1/2.

5. Thus, c (the car is behind door A/ Monty opens door B) = 1/2* 1/3 ÷ 1/2

6. Thus, the chance that the car is behind door A given the fact that Monty opens door B is 1/3.

7. Therefore, the chance that the car is behind door C is 2/3, and the switch strategy will be more successful in getting the car.

Again, imagine a classroom situation. The professor presents the Monty Hall Puzzler, and motivates the intuition (this is extremely easy to do!) that it is 50-50 whether we go for the stick strategy or for the switch strategy. The professor then gives the proof from Bayes's Theorem above. Now suppose that the students have learned Bayes's Theorem—they understand it, believe it to be true, and know why it is true. So the students have done more than memorize the theorem, they actually know that it is true. Moreover, the students understand and accept the proof above, and thus know that Bayes's Theorem implies that it is better to switch. Despite all this, it still seems possible for a student to refuse to accept the conclusion—the intuition that it is just 50-50 is too hard to shake off.[8]

To be sure, there may be a certain amount of cognitive dissonance on the part of the students in these counterexamples (although I doubt even this to be a necessary consequence) when they reflect on the proofs and their attitudes towards the various parts of the proofs. Such distress or confusion may even result from a subconscious desire to adhere to principle K4. Nevertheless, principle K4 has been shown to be false. It is not necessarily true that knowledge is closed under known implication. Even to show that it is contingently closed would require a demonstration that the student-type cases do not occur in the actual world. If the students were perfectly rational subjects they would accept the conclusions that they in fact reject. But nothing is more commonplace than the fact that the world is populated with less than perfectly rational people.

One move, taken by Graeme Forbes (Forbes 1984) and Steven Luper-Foy (Luper-Foy 1984), is to modify K4 in the following way in an attempt to preserve its truth.

K5: $\square \ \forall \ x \ \forall \ y \ ((Kx \ \& \ K(x \rightarrow y) \ \& \ By) \rightarrow Ky)$
Closure of knowledge under known implication plus belief in the consequent

In other words, if you know x, you know that x implies y, and *you believe that y*, then you know that y. K5 is the natural successor to K4, since it successfully rules out the cases of the Logic Student and the Monty Hall Puzzler. However, K5 also has serious problems. Consider this case. Jane has given some thought to the Monty Hall Puzzler, and has been exposed to the Bayes's Theorem proof in class. While she knows that Bayes's Theorem is true, and she knows that this entails that you should switch, she just cannot bring herself to accept the conclusion. Over lunch she discusses the puzzler with one of her friends, a notoriously unreliable reasoner who is unusually susceptible to specious arguments. This friend, Dick, tells Jane that the switch strategy is superior after all, perhaps of-

fering up a transparently unsound argument in defense of this view. Jane tends to have implicit faith in her friends, takes Dick at his word, and so comes to believe that the switch strategy is best. Surely Jane's belief is unjustified. She is not basing her belief on the legitimate but unintuitive argument from Bayes's Theorem, but on the illegitimate say-so of her unreliable friend. Or to modify the case, suppose that Jane is as above, but instead of having lunch with Dick, she has a dream in which Elvis comes to her and tells her that the switch strategy is best. Jane always believes any messages Elvis delivers in her dreams, and so forms the belief on the basis of that. I assume that all hands would agree that her belief would be an unjustified one. If we add the common and plausible assumption that justification is a necessary condition for knowledge, it follows that Jane does not know that the switch strategy is best. Thus, Jane knows Bayes's Theorem, knows that it entails that you should switch, believes that you should switch, and yet she does not know that you should switch. Hence K5 is false.

Forbes and Luper-Foy might reply at this point that K5 is meant to capture the case in which belief in y is arrived at by deduction from x and x implying y. The idea is that when you know x, know x implies y, and deduce through modus ponens (and subsequently believe) y, then you know that y. The revised principle is then something like this:

K6: $\square$ $\forall$ x $\forall$ y ((Kx & K(x $\to$ y) & By & By is based on Bx and B(x $\to$ y)) $\to$ Ky)
Closure of knowledge under known implication plus belief in the consequent based upon deduction via modus ponens[9]

Notice that the method of deduction is built into K6 precisely to prevent arriving at a belief in y due to faulty belief-acquiring methods. In other words, to save the closure principle a good epistemic method had to be built into it. So deduction was added. *Deduction* was specified as a method because of an *antecedent* commitment to it as a good epistemic method. So of course K6 will be true—it just repeats in a fancy way what was antecedently ordained, viz., that deduction is a good epistemic method. Thus, K6 does not seem like much of a discovery. It may be true, but it seems trivially so: the insistence on deduction, coupled with the presumption that deduction used under these circumstances cannot fail to yield knowledge, is sure to produce knowledge.

One might fairly object here that even if K6 is trivially true in some sense, it is still *true*, and so a good principle to add to our stock of epistemological verities. That is, one should drive a wedge between being logically or formally trivial, and being

philosophically trivial. K6 may be formally trivial, given certain assumptions about the semantics of "know," but this alone does not show the principle to be philosophically trivial. Mathematical theorems may be formally trivial, but this does not show that they are therefore worthless or mathematically insignificant.

This objection is correct. However, it is difficult to evaluate the philosophical significance of K6 in abstraction. A proper evaluation requires an extended excursion into each project—skepticism, foundationalism, what have you—laying claim to a K6-like principle. As stated in the introduction, one program of doxastic logic is to provide principles that allow us to derive what someone believes on the basis of other things they believe. K6 is plainly philosophically trivial for this agenda, since what amounts to justified belief in the consequent is built right into the antecedent. There is no *deduction* of belief or justified belief to be had here. I have no general argument that other enterprises relying on K6 are philosophically bankrupt. It does seem, though, that given the evidence that K6 is formally trivial, the purveyors of these enterprises are under some obligation to show how their particular dependence on K6 is not philosophically trivial as well.

I conclude that closure principles for knowledge K1-K5 are false, and that K6 is trivially true. Doxastic logicians might at this point abandon hope for the closure of knowledge, and instead focus on closure principles for justification alone. As we will see, several philosophers have discussed such principles. Before delving into them, there is an important point to note about justification. Some justification locutions imply belief, for example, 'B is a justified belief', 'S's belief that p is justified', and 'S justifiedly believes that p'. Some locutions do not, e.g., 'there is justification for S to believe that p', or 'S has a justification to believe that p'.[10] Let us call belief-implying justification 'strong justification', and justification that does not imply belief 'weak justification'. I will use 'Jx' to stand for strong justification. Thus, $\Box \, \forall \, x(Jx \to Bx)$. This is how justification has been treated previously in principle K3. However, in dealing with principles meant to show entailment relations solely concerning justification, it will be valuable to explore weak justification as well. I will use '$\mathcal{J}x$' to stand for weak justification. Consider then J1.

J1: $\Box \, \forall \, x \, \forall y \, ((Jx \, \& \, (x \to y)) \to Jy)$
    Closure of strong justification under implication

It will be noticed that J1 is even stronger than K1. Recall that K1 states that all logical implications of knowledge are themselves known. J1 is also a kind of omniscience principle, as it states that all of the logical implications of one's justified be-

liefs are themselves justified beliefs. As in the case of knowledge, it is extremely unlikely that anyone really believes the myriad things entailed by their justified beliefs. There is just too much to believe.[11] Let us turn then to the analogous principle for weak justification.

$\varnothing1$: $\Box$ $\forall$ x $\forall$y $((\varnothing x \,\&\, (x \to y)) \to \varnothing y)$
Closure of weak justification under implication

This principle is advanced by Peter Klein (Klein 1981, 41–81) and Robert Koons (Koons 1992). The problem is that it is difficult to evaluate $\varnothing1$ without having a better grasp on just what constitutes weak justification. An example Jonathan Kvanvig gives to illustrate justification that does not imply belief may be helpful here. Suppose that Joe is driving from Dallas to Houston at 45 MPH. Given that the speed limit is 55 MPH, we might say that Joe's driving 55 is justified, even though he is not driving that fast (Kvanvig 1992, 68). Likewise, S can be justified in believing x, even though S does not believe x. The idea seems to be that there is sufficient evidence available to Joe to justify driving 55 (a posted speed limit sign, say), or adequate evidence available to S to justifiably believe that x.[12] Thus, if Joe were to become aware of this evidence (he sees the speed limit sign) and act on the basis of it (speed up to 55), then this behavior would be justified. Likewise, if S were to recognize the evidence for x and form a belief that x on the basis of it, then S's belief in x would be justified. Thus, weak justification seems to be a kind of counterfactual or dispositional strong justification.

Under this interpretation $\varnothing1$ turns out to be false. To see this, suppose that S is weakly justified in believing x. That is, if S were to believe x, then S's belief would be a (strongly) justified one. Let x be "it is raining outside." If S were to believe x, she would acquire the belief by looking outside and seeing the storm, thus coming to have a (strongly) justified belief. Further, suppose that x implies y: S will not have to water her tomatoes tomorrow. Do these things imply $\varnothing y$; i.e., that if S were to believe that she will not have to water her tomatoes tomorrow, then this belief would be (strongly) justified? Note that nothing about $\varnothing1$ requires the counterfactual belief in y to depend upon a counterfactual belief in x. Nor is there anything that would hitch the counterfactual belief in y to recognizing that x→y. Thus, nothing prevents us from supposing that if S were to believe that she will not have to water her tomatoes tomorrow, she would come by this belief by asking her completely unreliable neighbor if the tomatoes need to be watered. In such a case, S is weakly justified in believing x, x→y, and yet y is not weakly justified for S; thus $\varnothing1$ is false.

One could counter that there are other ways of interpreting weak justification besides as counterfactual or subjunctive strong justification. We might explain it this way: a proposition x is weakly justified just in case it is a *good thing* to believe x. This is not to say that it *would* be a good thing if one *were* to believe it, or that one *would* be justified in doing so. Just that if believing a proposition x is a good thing, then x is weakly justified. We could then cash out "good thing" talk in terms of subjective probability.[13] Thus, $\wp$x means that there is a certain high probability that x is true, or that the probability of x is over some specified threshold. A subject is justified in believing x (even when the subject does not believe x) just in case x is highly probable. We might think of it like this: according to the traditional epistemic enterprise of trying to acquire true beliefs and avoid false ones, we will do a pretty good job if we believe only those things which are highly probable given our evidence. We need not insist on a subjective probability of 1—this is too much to ask. Lower thresholds—90% say—are enough to make us justified in believing.

However, on the analysis of weak justification as subjective probability, it is hardly a straightforward matter that $\wp$1 is true. If we specify x to be an atomic proposition, then $\Box$ $\forall$ x $\forall$y (($\wp$x & (x → y)) → $\wp$y) turns out to be a truth of probability theory. However, suppose that x is a conjunctive compound composed of at least two atomic propositions. Richard Foley, for example, has argued that it is possible to have justified inconsistent beliefs—and so presumably there are inconsistent sets of propositions such that each element of the set is highly probable given the evidence available to the subject (Foley 1979). Given this, it seems possible that an inconsistent conjunctive compound can be weakly justified. For example, suppose that we set our threshold at .7, and x is the inconsistent compound composed of p & q. If the evidence in favor of p makes it 90% probable, and the available evidence for q makes it 80% probable, then x is 72% probable and so weakly justified.[14] Therefore, under $\wp$1 everything is weakly justified for an subject for whom x is weakly justified, since an inconsistency entails everything. Since plainly not every proposition is weakly justified (even for a subject for whom a conjunction of inconsistent propositions is weakly justified), it seems that a defender of $\wp$1 must deny the possibility of a weakly justified inconsistent compound.

There are obvious ways to try to modify J1 and $\wp$1. Here is the first:

J2: $\Box$ $\forall$ x $\forall$y ((Jx & B(x → y)) → Jy)
Closure of strong justification under believed implication

This is not much of an improvement over J1, as the Logic Student and Monty Hall Puzzler cases serve as counterexamples to it. Again, belief in y cannot be guaranteed in the consequent. What about $\mathcal{J}2$?

$\mathcal{J}2$: $\Box \ \forall \ x \ \forall y \ ((\mathcal{J}x \ \& \ B(x \to y)) \to \mathcal{J}y)$
Closure of weak justification under believed implication

This may even be worse than J1. Belief that x implies y is insufficient for transmitting justification. After all, it might be the case that if S were to believe x then he would be justified in doing so, yet S's actual belief that x implies y is wildly ungrounded. This provides no reason for thinking that S is weakly justified in believing y. To point this up, suppose under the counterfactual interpretation that if S were to believe that the planets were aligned in such-and-so a way, then he would be justified in so believing. Suppose that S superstitiously believes that if the planets are aligned in this way, then doomsday is near. We should in no way conclude from this that if S were to believe that doomsday is near, that he would be justified in doing so. The probabilistic interpretation is no better—the fact that the subject believes x to imply y tells us nothing about the probability of y, unless we also know something about the probability of x implying y.

What about the next logical amendment?

$\mathcal{J}3$: $\Box \ \forall \ x \ \forall y \ ((\mathcal{J}x \ \& \ \mathcal{J}(x \to y)) \to \mathcal{J}y)$
Closure of strong justification under strongly justified implication

Roderick Chisholm (Chisholm 1989, esp. 54; cf. Chisholm 1976, Appendix D) and James Van Cleve (Van Cleve 1979) express sympathy towards, if not downright endorsement of, this principle. Nevertheless, it too is false, and falls directly to the Logic Student and Monty Hall Puzzler counterexamples. Someone could justifiably believe the premises of the Bayes's Theorem proof, justifiably believe that they imply that it is better to switch, and still fail to believe (and thus fail to justifiably believe) the conclusion. Let us turn to its analogue $\mathcal{J}3$.

$\mathcal{J}3$: $\Box \ \forall \ x \ \forall y \ ((\mathcal{J}x \ \& \ \mathcal{J}(x \to y)) \to \mathcal{J}y)$
Closure of weak justification under weakly justified implication

By now it should be fairly obvious how to construct a counterexample to the counterfactual interpretation of this principle. The case of the tomato grower considered previously should do it. S is weakly justified in believing that it is raining outside

(if she were to believe it, then her belief would be justified). Suppose that S is also weakly justified in believing that if it is raining outside, then she will not need to water her tomatoes tomorrow (if she were to believe this conditional, her belief would be justified). Nevertheless, S is not weakly justified in believing that she will not need to water her tomatoes tomorrow, since if she were to acquire this belief, she would do so by asking her unreliable neighbor if the tomatoes will need watering.

What of the probability interpretation of $\phi 3$? This is subject to Henry Kyburg-type arguments against "conjunctivitis" (Kyburg 1970). Suppose that the threshold for weak justification is a probability of 70%. Now assume that there is an 80% chance that x is true, and an 80% chance that x implies y is true. The chance of y conditional upon the chance of x and the chance of x implying y (as specified by $\phi 3$) is no greater than the product of the chance of x and the chance of x implying y, i.e., 64%. As 64% is less than our threshold of 70%, y is not weakly justified, even though x and x→y are, and so $\phi 3$ is false.[15]

There are no doubt other ways to cash out weak justification besides subjunctive strong justification or subjective probability. Perhaps there are closure principles that hold for these other analyses—it is difficult to say a priori. Of course, one might just declare that whatever weak justification is, it is the kind of thing that supports a closure principle and avoids the counterexamples given above. Such a move would be rather like defining knowledge as justified true belief plus whatever avoids the Gettier cases. While it gets you what you want, it is rather ad hoc and unsatisfying.

I have nothing to say here about the relationship between the failure (or triviality) of closure principles and the analysis of knowledge as justified true belief, nor will I say anything about the prospects for skepticism in light of these results. I note in passing the challenge laid at the door of foundationalists: if our mediate knowledge is grounded in or based upon an immediate foundation, none of the dozen closure principles discussed in this essay adequately characterizes a transmission principle.[16] A quite striking result to be highlighted is the blow struck to doxastic logic. In one sense, doxastic logic is nothing other than the time-honored affair of conceptual analysis, showing the relations among our various epistemic concepts. The results of this paper show that many of these relations often thought to hold do not in fact do so. Of course, this may only send analysts back to the drawing board to try to find new principles that work. And, probably, some improvements will be forthcoming. However, not in the offing are non-trivial necessary truths that allow us to conclude what someone believes or knows on the basis of other things that

they believe or know. This is a lesson of the Logic Student and Monty Hall Puzzler cases. Insofar as doxastic logic aims at such results, it is a quixotic enterprise that has no payoff for epistemology. There may be contingent principles that accurately tell us what someone is likely to believe given other things that they believe, but such principles will not be discovered through logic or conceptual analysis. Such truths will only be discovered through empirical study, and thus the promotion of these contingent claims should be left to the naturalizers among us.[17]

NOTES

[1]   It has been suggested to me that computers are not fictional, and are epistemically interesting. However, the attribution of intentional states to computers is so controversial that saying anything about the truth of closure principles for them requires another paper.

[2]   Compare Carnap 1959, especially page 68 on logical vs. grammatical syntax, with Gamut 1991, 72–73: "what we are trying to describe is not how we should speak if we want to please philosophers but how we do in fact speak."

[3]   "→" in this and all following principles is to be understood as material implication.

[4]   It should be noted that in Footnote 1 of his paper Vogel claims that K4 needs a further refinement in order to rule out the case in which someone might know p, know p implies q, and fails to "put these things together" and so does not infer q. It may seem that the counterexamples to be considered later are of the sort Vogel wishes to rule out. They are not. In the upcoming counterexamples, the subjects are occurrently attending to both their knowledge that p and their knowledge that p implies q. They *are* putting things together. Nevertheless, they still cannot get to q.

[5]   Additional opposition to McGinn is available to those who think that knowledge is essentially contextual. If it is, then regular knowledge claims (e.g., I know that I am in Pennsylvania) can be true in ordinary contexts C, but false in other, stricter, contexts C*. When the skeptic raises brain-in-vat possibilities, he likewise raises the standards for knowing, and so shifts to C*. McGinn's criticism of K4 relies on illicitly moving between contexts. A genuine counterexample to K4 would show that one can know p in context C, know p implies q in C, and fail to know q in C. McGinn shows that you fail to know q in C*.

[6]   From Quine 1982, 184. The following proof is also Quine's.

[7]   Thanks to James Dreier for supplying this proof.

[8]   In fact, anecdotal evidence suggests that this actually happens. Students turn out to be more easily persuaded that they should switch by playing Monty Hall-type games than by the formal proof.

[9]   Compare Edmund Gettier: "for any proposition P, if S is justified in believing P, and P entails Q, and S deduces Q from P and accepts Q as a result of this deduction, then S is justified in believing Q" (Gettier 1963, 121).

[10]   For a detailed discussion of different justification locutions, see Kvanvig 1992, ch. 4. Also see Alston 1989.

[11]   Ernest Sosa also rejects this principle. See Sosa 1991, 24.

Steven D. Hales

¹² This seems to be the tack of Klein. See his discussion of "available evidence" on 45–48.
¹³ This seems to be the approach of Koons.
¹⁴ Note this last will follow directly from Foley's inconsistent sets plus the closure of weak justification under conjunction. Since this principle is probably false—as Foley among others argues—an argument was needed to show that an inconsistent conjunctive compound can be weakly justified.
¹⁵ Koons defends ℘3, but wisely restricts it to cases where the threshold for justification is probability 1. See Koons 1992, 16.
¹⁶ I am assuming that foundationalists will not be satisfied with the strong Cartesian sentiments expressed in K6. See Van Cleve 1979.
¹⁷ Thanks to Robert Almeder, John Carroll, Richard Feldman, Richard Kirkham, and Steven Rieber for criticisms of an earlier draft. Thanks also to Richard Feldman for allowing me to see his unpublished work.

REFERENCES

Alston, W. P. 1989. *Epistemic Justification*. Ithaca, New York: Cornell University Press.
Brueckner, A. 1985a. Skepticism and epistemic closure. *Philosophical Topics* 13 (3): 89–117.
―― 1985b. Transmission for knowledge not established. *Philosophical Quarterly* 35 (139): 193–195.
―― 1994. The structure of the skeptical argument. *Philosophy and Phenomenological Research* 54 (4): 827–835.
Carnap, R. 1959. The elimination of metaphysics through the logical analysis of language. In *Logical Positivism*, ed. A. J. Ayer, 60–81. New York: Macmillan.
Chisholm, R. M. 1976. *Person and Object*. La Salle, IL: Open Court.
―― 1989. *Theory of Knowledge*. Englewood Cliffs, NJ: Prentice Hall, 3rd ed.
Dretske, F. 1970. Epistemic operators. *Journal of Philosophy* 67 (24): 1007–1023.
Feldman, R. 1994. In defense of closure. Presented at the Pacific Division meetings of the American Philosophical Association.
Foley, R. 1979. Justified inconsistent beliefs. *American Philosophical Quarterly* 16 (4): 247–257.
Forbes, G. 1984. Nozick on skepticism. *Philosophical Quarterly* 34 (134): 43–52.
Gamut, L. T. F. 1991. *Logic, Language, and Meaning, Volume II: Intensional Logic and Logical Grammar*. Chicago: University of Chicago Press.
Gettier, E. L. 1963. Is justified true belief knowledge? *Analysis* 23: 121–123.
Hintikka, J. 1962. *Knowledge and Belief*. Ithaca, New York: Cornell University Press.
Klein, P. D. 1981. *Certainty: A Refutation of Skepticism*. Minneapolis: University of Minnesota Press.
Koons, R. C. 1992. *Paradoxes of Belief and Strategic Rationality*. Cambridge: Cambridge University Press.
Kvanvig, J. L. 1992. *The Intellectual Virtues and the Life of the Mind*. Savage, MD: Rowman and Littlefield.

Kyburg, H. E. Jr. 1970. Conjunctivitis. In *Induction, Acceptance, and Rational Belief*, ed. M. Swain, 55–82. Dordrecht, Holland: Reidel.

Luper-Foy, S. 1984. What skeptics don't know refutes them. *Pacific Philosophical Quarterly* 65: 86–96.

McGinn, C. 1984. The concept of knowledge. *Midwest Studies in Philosophy* 9: 529–554.

Nozick, R. 1981. *Philosophical Explanations*. Cambridge: Harvard University Press.

Quine, W. V. O. 1982. *Methods of Logic*. Cambridge: Harvard University Press, 4th ed.

Sosa, E. 1991. *Knowledge in Perspective*. Cambridge: Cambridge University Press.

Stine, G. C. 1976. Skepticism, relevant alternatives, and deductive closure. *Philosophical Studies* 29: 249–261.

Van Cleve, J. 1979. Foundationalism, epistemic principles, and the Cartesian circle. *Philosophical Review* 88: 55–91.

Vogel, J. 1990. Are there counterexamples to the closure principle? In *Doubting*, eds. M. D. Roth and G. Ross, 13–27. Dordrecht, Holland: Kluwer.