



# Closure Principles

Jonathan L. Kvanvig

University of Missouri

---

## Abstract

A dispute in epistemology has arisen over whether some class of things epistemic (things known or justified, for example) is closed under some operation involving the notion of what follows deductively from members of this class. Very few philosophers these days believe that if you know that  $p$ , and  $p$  entails  $q$ , then you know that  $q$ . But many philosophers think that something weaker holds, for instance that if you know that  $p$ , and  $p$  entails  $q$ , then you are in a position to know that  $p$ , or if you know that  $p$  and you competently deduce  $p$  from  $q$ , then you know that  $q$ . However there are some considerations tracing back to Wittgenstein's *On Certainty* and some early work by Fred Dretske that suggests even these weaker principles are false.

---

One of the most significant disputes in epistemology over the last forty years or so has been about closure principles. Here are two paradigm examples of closure principles in epistemology:

- (1) If you are justified in believing  $P$ , and  $P$  entails  $Q$ , then you are justified in believing  $Q$ .
- (2) If you know  $P$ , and  $P$  entails  $Q$ , then you know  $Q$ .

We can begin by explaining what makes these principles count as closure principles and the nature of the controversy about them.

A *closure principle* is a principle that claims that a certain category of object (typically a set) is closed relative to some function or operation or rule, in the sense that performing that operation on any member of the set always leads us to something already in the set. Consider, for example, the natural numbers and the operation of addition. Speaking metaphorically, we might describe the door to this set as being closed under this operation, since whenever we add two natural numbers, we are led to something already in the set. By contrast, when a set is not closed under a particular operation, the door stands open and non-members may be encountered when the operation is performed. So, for example, the natural numbers are not closed under subtraction: when we subtract one natural number from another, the result may lie outside the set (it might be a negative number).

An *epistemic principle* is a conditional that involves epistemic concepts (such as justification, knowledge, being in a position to know, being warranted,

having positive epistemic status, being certain, and the like). Some examples of epistemic principles include: if you are justified in believing a claim, then you are in a position to know it; if a claim is certain for you, then you know it; if you know a claim, then it has positive epistemic status for you; etc. An *epistemic closure principle* is an epistemic principle that identifies a group or set falling under a certain epistemic concept, and says that under a certain operation or function or rule, that set is closed.

A dispute in epistemology has arisen where the operation in question has something to do with deductive consequences. The question is whether some class of things epistemic (things known or justified, for example) is closed under some operation involving the notion of what follows deductively from members of this class. As a result, a paradigm case of an epistemic closure principle is one that says that the set of things you know is closed under entailment. In ordinary English, this claim can be expressed as claim (2) above: that if you know  $p$ , and  $p$  entails  $q$ , then you know  $q$ .

Closure principles such as (2) are easily shown to be false, however.<sup>1</sup> People do not believe all the logical consequences of their beliefs, and hence do not know all the logical consequences of their beliefs. Moreover, mathematical and logical truths follow from any set of statements whatsoever, so if we know or are justified in believing anything, it follows from this fact and the above principles that we know and are justified in believing all logical and mathematical truths. It seems obvious that we are not logically omniscient in this way, and it is equally obvious that we are sometimes justified in believing the denial of a logical truth. These are just some of the ways in which principles (1) and (2) above are defective.

The primary controversy, however, is not about (1) and (2), but rather about some of their close cousins. I'll begin by focusing on principles about knowledge and broaden the discussion later. It is a well-known fact about deduction that it is a way of extending our knowledge. Sherlock Holmes extends his knowledge in this way in solving crimes, as do the rest of us when we think through the consequences of various plans we might adopt. Since it seems obvious that our knowledge can be extended in this way, it seems that we should be able to formulate some close variation on (1) or (2) that captures this fact. For example, instead of knowing all the consequences of one's knowledge, perhaps a correct principle stresses that one needs to recognize the implications of what one knows in order for one to know or justifiably believe what is implied.

Many philosophers have argued, however, that no close cousin of the above principles could be true. The history of arguments against closure principles traces opposition to two sources. The first source is in the discussion of lottery and preface paradoxes. In the lottery paradox, the problem arises if it is granted that you can know that your ticket will lose. Since the same basis for such knowledge would apply to all tickets, it appears that one could thereby infer that no ticket will win, contrary to the known fact that some ticket will win (in one ordinary kind of lottery). In the preface paradox, the author of a

book writes in the preface that though he or she has been conscientious in eliminating error, mistakes are sure to remain in the body of the work. If the author has good reasons for believing everything in the book and good reasons for believing the preface statement, it appears that good reasons can be found for believing an explicit contradiction, since the preface statement is inconsistent with the claim that the contents of the book are all true.<sup>2</sup>

The other primary source of discontent with closure principles comes from epistemologies that understand knowledge in terms that advert, either implicitly or explicitly, to merely possible worlds. Such theories are *modalist epistemologies*: they offer accounts of knowledge that explain knowledge modally, in terms of what is true or false in alternative possible situations. One such theory of knowledge is *infallibilism*, according to which knowledge requires that error be impossible (given the grounds for the belief that obtain). Though this extreme version of modalist epistemology has never been used to question closure, there are weaker versions of modalist epistemology that have been so used – versions that were developed out of a dissatisfaction with infallibilism. On the infallibilist view, one is in a position to know something only if one's grounds for belief make error impossible; they allow one to rule out every possible alternative to what one actually believes. Perhaps, though, that is too strong; perhaps not every alternative needs to be ruled out, since some alternative possibilities may be so remote as to be irrelevant in a given context. Such a theory claims that knowledge requires ruling out relevant alternatives, but it refuses to claim that every possible alternative is a relevant alternative in every context.

Alvin Goldman's "Discrimination and Perceptual Knowledge" (1976) gives classic expression to the relevant alternatives approach to knowledge; according to this view knowing doesn't require ruling out *all* possible alternatives to what one presently believes, but only the *relevant* alternatives, where the notion of relevance is explained differently by each different version of the view. To get from here to a denial of closure, one need only add that, where  $p$  entails  $q$ , the alternatives to  $p$  that are relevant need not be the same as the alternatives to  $q$ . The first explicit denial of closure based on this idea is found at the end of Fred Dretske's seminal paper "Epistemic Operators" (1970) where, in the context of a more general discussion of epistemic operators, Dretske maintains that the knowledge operator preserves closure for some logical connectives but not for all.

To use Dretske's famous example, suppose you are visiting the zoo and notice a zebra. You know that it is a zebra because you can rule out relevant alternatives to your belief that it is a zebra: you can rule out that it is an elephant, a lion, a tiger, and all of the other animals usually found in zoos. Your belief entails the falsity of all of these alternatives, but it also entails that the animal is not a cleverly disguised mule. Once clever disguises are introduced into the discussion, however, the capacity to rule out relevant alternatives disappears. You couldn't tell the difference, for example, between a cleverly disguised mule and a cleverly disguised donkey, presumably a

relevant alternative to the mule hypothesis. So you don't know that it is not a cleverly disguised mule. On the relevant alternatives approach to knowledge, however, there is no need to worry. You know that the animal is a zebra even though you can't rule out the alternative that it is not a cleverly disguised mule. You don't know that it is not a cleverly disguised mule, but that doesn't undermine your knowledge that it is a zebra, because even though the claim that it is a cleverly disguised mule is an alternative hypothesis to the one you accept, it is not a *relevant* alternative.

Viewed in this way, the relevant alternatives theory has a precursor in Wittgenstein's *On Certainty* (1969).<sup>3</sup> Wittgenstein, in his opposition to Moore, distinguishes between what we know and the grounds of what we know, where the grounds themselves need not count as knowledge. In applying this distinction, Wittgenstein returns repeatedly to claims of ordinary life for the first category (what we know) and philosophical claims for the second category (the grounds that we take for granted, but do not know). Moore notoriously claimed, however, that he could prove and hence know the truth of philosophical claims such as the claim that there is an external world on the basis of knowing claims of ordinary life, such as the claim that here is one hand and here is another. Moorean anti-skepticism is rooted in the idea that we can extend our knowledge by deducing philosophical claims from ordinary items of knowledge, and Wittgenstein's opposition to Moore on this point requires a denial of closure of the same sort defended by the relevant alternatives theorists who explicitly deny closure.

Besides by appeal to the relevant alternatives theory, another way to question closure is through epistemologies that understand knowledge in terms of counterfactual connections between a belief and its truth. The counterfactual that has been used to question closure is termed the *sensitivity* condition: that if a belief is to count as knowledge, then it must be the case that if the content of the belief in question were false, the belief would not be held. Robert Nozick's theory of knowledge (1981) employs just such a sensitivity requirement as a distinctive element differentiating true belief from knowledge. On the usual semantics for counterfactuals,<sup>4</sup> we interpret a counterfactual by looking for close worlds in which the antecedent is true and check to see whether the consequent is true as well, so such theories count as versions of modalist epistemology as well. Given this understanding of counterfactuals, even if  $p$  entails  $q$ , the close  $\sim p$  worlds and the close  $\sim q$  worlds need not be the same, and it is this feature of the account that undermines closure. Nozick writes:

So, we must adjust to the fact that sometimes we will know conjunctions without knowing each of the conjuncts. Indeed, we already have adjusted. Let  $p$  be the statement that I am in Emerson Hall, not-SK be the one that I am not on Alpha Centauri floating in that tank; since  $p$  entails not-SK,  $p$  is (necessarily) equivalent to  $p \ \& \ \text{not-SK}$ . I know that  $p$ , yet I do not know that not-SK. (228)

The particular deductive connection that Nozick here denies is closed for knowledge is that described by the conjunction elimination rule (the rule

according to which A, as well as B, follows from the conjunction of A and B), but Nozick's theory provides a basis for abandoning a variety of deductive connections as knowledge-transmitting.

Once we appreciate the sources of discontent about closure, we can see that the debate about closure is not simply a debate about whether there is a true closure principle. The sources of discontent threaten not only epistemic closure principles, but other, non-closure epistemic principles as well. Consider, for example, Dretske's painted mule example mentioned earlier. If he is right, knowing that the object is a zebra is relatively easy but knowing that the animal is not a painted mule is hard. Such an assessment threatens closure, but it also threatens the idea that if you know a claim is true and recognize its implications, you might still lack the epistemic grounds needed for knowledge. In the example above, according to Nozick, he knows he is in Emerson Hall, but lacks the requisite grounds for knowledge that he's not on Alpha Centauri floating in a tank. So not only does the example threaten this closure principle:

(3) If you know P, and know that P entails Q, then you know Q;

it also threatens a much more plausible epistemic principle (which is, strictly speaking, not a closure principle):

(4) If you know P, and know that P entails Q, then you are in a position to know Q,

where we understand "in a position to know" in terms of possessing the epistemic grounds needed for knowledge.

It is a legitimate complaint against the present state of the literature that this latter principle is also called a closure principle. It is clearly not a closure principle, since, as pointed out earlier, a closure principle in this context will begin with the set of things known and identify some function or operation on members of that set that yields only results that are already in the initial set of things known. We should not make too much of this point, however, for the sloppiness results from some implicit awareness that the attacks on closure, if successful, have implications beyond closure itself. If this broader range of principles contains some easily defensible principles, then the arguments against closure should be rejected.

Even so, the debate can still be cast as focusing on closure principles themselves, since the additional elements needed to move from, say, being in a position to know to actually knowing can be built into the operation involved in some closure principle. For example, return to the point noted above that a primary motivation for endorsing closure under implication for knowledge is that deduction is a way of extending knowledge. One way to describe this idea is by limiting the operation in question to known entailment, and then claiming that under this operation, knowing P implies being in a position to know P. Another option, however, is to replace the operation in question with a related one – say, the operation of having

competently deduced one claim from another. Broadening the operation in this way allows formulating a closure principle along the lines proposed by Timothy Williamson (2000):

(5) If you know P, and competently deduce Q from P, then you know Q.

Notice that this principle is a closure principle, and not simply an epistemic principle somehow involving the notion of entailment. It takes a class of things epistemic, in this case the class of things known, and says that anything you've competently deduced from a member of this class is already a member of the class. Hence, by broadening the kinds of operations allowed on a class of things epistemic, the non-closure principles involving entailment can typically be recast in terms of some closure principle falling under a broader operation than mere entailment or known entailment (I use the term "broadening" here because the notion of competent deduction is not a mathematical operation in the strict sense). For this reason, the disputes about related epistemic principles can be pursued focusing only on closure principles involving broader operations that nonetheless involve the notion of entailment of deductive consequences, such as the Williamsonian principle articulated above.

There are various reasons to hold that this principle (5), even if on the right track, is in need of some refinement. For one thing, a person might competently deduce Q from P without coming to believe Q: as Gilbert Harman has reminded us for decades, when you deduce consequences of what you believe, sometimes you should add a new belief and sometimes you should abandon old beliefs.<sup>5</sup> In addition, the process of deduction might itself generate new information independent of the information involved in the deduction itself, and this new information may itself be relevant to whether one comes to know Q by competently deducing it in this way.

In light of these and other considerations, John Hawthorne, while recommending the basic approach adopted by Williamson, has suggested the following refinement of this view (29):

(6) If one knows P and competently deduces Q from P, thereby coming to believe Q, while retaining one's knowledge that P, one comes to know that Q.

This approach needs one further small emendation to account for the possibility that one learns of a (misleading) defeater<sup>6</sup> for Q while competently deducing it from P. In such a case, one will not be able to come to know Q in this way. The emendation needed gives us:

(7) If one knows P and competently deduces Q from P, thereby coming to believe Q, while retaining one's knowledge that P and learning of no undefeated defeater for Q in the process, one comes to know that Q.

There is little reason to try to avoid this emendation. Rarely, in the process of deducing one claim from another, do we learn nothing else. We may be

watching television, looking at the clouds in the sky, or any number of other things. In doing these other things, we may, for example, overhear a misleading defeater for what we are inferring, but simply ignore its implications for whether we should believe Q. So we come to believe Q by competently deducing it from P. But, if the defeater is of the right sort,<sup>7</sup> we won't know Q, because of the defeater. (As it happens, it is a misleading defeater, but we are unaware of that part of the story.) So, for example in the case of deduction, you might overhear two logicians trying to rob you of knowledge, saying to each other that the inference from P to Q is (subtly) invalid. You hear what they say and ignore it, forming the belief that Q on the basis of competent deduction from P. But you don't know Q.<sup>8</sup>

So Hawthorne's principle needs fixing. The small emendation suggested fixes this problem, but questions about the truth of the principle remain. Some of these questions involve details about the principle. For example, the principle employs the notion of competent deduction, but in ordinary parlance, a person can deduce things that don't follow and can do so competently. The truth of the principle requires, however, that the deduction be correct; so, we'll need to understand competence so that it involves the notion of correctness. More information is needed as well about the notion of competence itself, since presumably there will be vague cases where it is not clear whether a particular correct deduction counts as a competent one.

Besides these matters of detail, there is the matter of defense of the principle. A primary reason for endorsing such a closure principle derives from the truism that deduction is a way of extending our knowledge. That reason, however, is not conclusive, since that truism would be sustained even if sometimes deduction failed to extend knowledge. In order to derive a principle having the exceptionless character of a closure principle, something more will be needed. Here it is tempting to appeal to features of conversational propriety. John Hawthorne gives just such an argument. He writes,

I ask S whether she agrees that P. She asserts that she does: "Yes," she says. I then ask S whether she realizes that Q follows from P. "Yes," she says. I then ask her whether she agrees that Q. "I'm not agreeing to that," she says. I ask her whether she now wishes to retract her earlier claims. "Oh no," she says. "I'm sticking by my claim that P and my claim that P entails Q. I'm just not willing to claim that Q." Our interlocutor now resembles perfectly Lewis Carroll's Tortoise, that familiar object of ridicule who was perfectly willing to accept the premises of a *modus ponens* argument but was unwilling to accept the conclusion. (32)

Hawthorne takes this example as illustrating the difficulties that attach to denying closure principles about knowledge in light of the knowledge norm of assertion, according to which you aren't supposed to assert what you don't know to be true. Something like this norm is needed for the impropriety in the above conversation to count as evidence for a closure principle about knowledge (though it appears that what is needed is

something different from the knowledge norm, something to the effect that what your committed to asserting is plausibly taken to be something you know). In any case, the example is intended to show that there are important connections between what we say or are committed to endorsing verbally and important epistemic conditions such as knowledge.

This argument from conversational propriety constitutes a challenge to those wishing to deny closure principles. The challenge issued is to insist that opponents of closure find a way to fill in the details of the example in a way that allows S not to be a ridiculous assessor, but a sensible one. If the challenge cannot be met, the argument from conversational propriety provides a strong basis for endorsing a closure principle about knowledge. If the challenge can be met, however, the argument will not be compelling, and there is reason to think that the challenge can be met.

The recipe for such an expansion of the case relies on the Problem of Easy Knowledge articulated by Stewart Cohen (2002). A simple example of this phenomenon is as follows. Suppose your son asks you whether there is a red object on the table, and you answer “yes.” Suppose he is suspicious of the claim, and wonders whether the lighting makes it appear red when it isn’t, and asks you about the possibility. If a theory of knowledge allows you to infer and thereby come to know that the lighting conditions aren’t worrisome from your (purported) knowledge that the object is red, you have acquired knowledge that the lighting conditions aren’t worrisome too easily – you have become an illustration of the problem of easy knowledge. As background for the expanded conversation, suppose S thinks about epistemology on occasion, and thinks that some things can’t be known on the basis of other things. She isn’t sure whether there is a way of ordering claims, even a partial ordering, so that we can always say what has to be known on the basis of what else. So she sticks with the negative characterization of epistemic priority just given: even if we can’t say what the order is exactly, we can say that some things can’t be known on the basis of some other things.

S is talking with Fred Dretske, in the presence of John Hawthorne:

Dretske says to her, “Is that a zebra?” and she says, “Yes.” Dretske says, “Do you realize that it can’t then be a disguised mule?” And she says, “Yes,” again. Dretske smiles, but Hawthorne is perturbed. Hawthorne gets her attention by tapping her on the shoulder: “so it’s not a disguised mule, right?” She says, “I’m not comfortable saying that yet. I can see it’s a zebra, and that this implies that it’s not a disguised mule. But I can’t learn that it’s not a disguised mule *like that*. And now that Fred points out the implication to me, I’m unsure of how or whether I know that it’s not a disguised mule.” Hawthorne then says, “So, I take it, you’ll want to retract one of your earlier claims.” “No,” she says, “I’m not questioning whether it’s a zebra – it is, I saw it – but I’m not sure I should be claiming that it’s not a disguised mule, since I can’t know that claim by inferring it from the zebra claim, and without knowing it in that way, I’m really not sure how I know it at all. So I think for now I should just remain silent on the point until I get this figured out a bit better, or until I go check to make sure



it's not a disguised mule." Hawthorne presses a bit: "You do agree that it follows from its being a zebra that it's a non-mule, right?" "Yes, but I hesitate, not because it might not be a non-mule, but because I'm not sure I'm entitled to be saying anything about it at this point."

The point of expanding the story in this way is this. S is not a ridiculous assertor in this story. She is confused, and there are various possibilities to explain how she can sensibly find herself in this predicament. One is Dretske's explanation that closure is false; another is that closure is true, and she's presupposing the idea that you shouldn't assert a claim when you're not sure that, or how, you know it to be true. Whatever the explanation, however, one can't accuse her of being like Lewis Carroll's tortoise, and so one can't infer from the expanded story that denials of closure are somehow incompatible with facts about conversational propriety.

Dretske's own hypothesis (2005) as to why closure fails appeals to the notion of a *heavyweight proposition*, a paradigm example of which is the painted mule claim used above. Dretske has no general account of what makes a proposition heavyweight, but intends the class to include philosophical propositions and skeptical hypotheses of the sort that underlie philosophical skepticism. This strategy is instructive, since the obvious strategy for denying closure while acknowledging the truism that deduction is a way of extending our knowledge is to restrict such extension to certain kinds of propositions. Any principle that carves the universe of propositions into those that are suitable instantiations of that principle and those that are not is clearly not a closure principle.

But Dretske's hypothesis fails as well, for if we change the context, we can find acceptable instances of precisely those inferences ruled out by that hypothesis. Suppose one is casually surveying one's zoo experience after the fact. One can reflect on the carefree afternoon, in conversation: "We saw lion, tigers, bears; mountain goats, alligators, bison; what else?" "Oh, lots of things – birds, fish, and snakes." "Zoos are such a – oh, how to put it? – *natural* experience. Well, except for the cages and all that! But what I mean is that what we experienced today wasn't a human invention, it wasn't an artifact. We got to see the real thing. We weren't looking at *pictures* of animals on TV; we weren't seeing fancy *robotics* designed to mimic real animals; we weren't seeing *people* dressed in *costumes* to look like the animals. Nature is beautiful and seeing it is so refreshing!" In such a conversation, there is no hesitation experienced anywhere in the conversation, and yet precisely the inferences barred by Dretske's hypothesis occur.

We thus have two conflicts to explore in order to resolve the question about closure. The first conflict arises from the argument from conversational propriety. In many, perhaps even most, contexts, failure to commit to acknowledged known consequences of what one asserts makes one appear ridiculous; but this feature is not always present, because of the problem of easy knowledge. Moreover, in many contexts, the problem of easy knowledge prevents inferences of the sort Dretske's hypothesis bars. But

again, that feature is not always present. So the challenge for defenders of closure is to give an argument for thinking that their favored principle is true, and they have not successfully discharged that responsibility to this point. For deniers of closure, the burden is to declare where the capacity to extend knowledge by deduction ends by delimiting the kinds of propositions we cannot learn simply by competently deducing them from what we already know.

### Notes

<sup>1</sup> At least under most views of belief. For a defense of such a closure principle in the context of a more general account of belief, see Stalnaker (1984).

<sup>2</sup> The lottery paradox was brought to the attention of the philosophical community by Henry Kyburg, who also argued that closure principles about conjunction must be denied in order to avoid the paradox (e.g., that if you know that your ticket will lose and know that your friend's ticket will lose, it doesn't follow that you know that both tickets will lose even if you recognize that this conjunction follows from your individual items of knowledge). The presentation appears in *Probability and the Logic of Rational Belief* in 1961, though the idea of denying closure is contained in his paper "Probability and Randomness" (1963), which was presented in 1959 at a meeting of the Association for Symbolic Logic. (Thanks to Greg Wheeler for pointing out to me the background of Kyburg's denial of closure in the context of the paradox.)

An earlier hint of such a position can be found in the work of F.P. Ramsey from 1929, who claimed that "we cannot without self-contradiction say p and q and r and . . . and one of p, q, r, . . . is false. . . . But we can be nearly certain that one is false and yet nearly certain of each; but p, q, r, are then infected with doubt" (110–11). (Thanks to Horacio Arlo-Costa for pointing out this source to me.)

Though D. C. Makinson is generally credited with discovering the preface paradox in 1965, it is clear that Ramsey's thought is an important predecessor of the version of this paradox, sometimes called the Fallibility Paradox, where the book is the set of one's beliefs and one's sense of one's own fallibility is contained in the preface to this book. The element needed for this version of the paradox, and which Ramsey does not mention, is the idea that one can be justified in believing claims that are only "nearly certain" for one. For a thorough treatment of these issues, see David Christensen (2004).

<sup>3</sup> Thanks to Peter Baumann for bringing to my attention the relevance of Wittgenstein to the issue of closure.

<sup>4</sup> Developed in detail in Lewis (1973).

<sup>5</sup> See, for example, Harman (1989).

<sup>6</sup> A *defeater* of the epistemic support relation between e and p is some further information d which, in conjunction with e, does not epistemically support p. For example, visual information of a certain sort supports the claim that there is a red object in front of one, but if one also knows that there is colored light shining on the object, this further information, together with the visual appearance, does not support this claim. The notion of a misleading defeater is a technical term in the literature on the Gettier problem, and it refers to defeaters of which one is unaware but which, unlike ordinary defeaters, do not imply that one does not have knowledge.

<sup>7</sup> Knowledge is undermined by *rebutting defeaters*: defeaters that provide evidence against the claim in question. For it is plausible to think that if E is evidence against Q, and P implies Q, then E is evidence against P as well (because it is evidence for  $\sim Q$  and hence evidence for  $\sim P$  because  $\sim Q$  implies  $\sim P$ ). (It is worth noting here that this claim about rebutting defeaters appears to depend on a closure principle about evidence, to the effect that if E is evidence for P, then it is evidence for anything that P entails as well.) The same cannot be said for *undercutting defeaters*: defeaters that threaten the connection between our evidence and what we believe. (Numerous examples of these sorts of cases can be found in the Gettier-inspired literature; the history of this literature is well-chronicled in Shope (1983).)

<sup>8</sup> Notice that we could avoid this counterexample if we tied the notion of competent deduction to an earlier closure principle. After noticing that the defeater in question is also a defeater for the

claim that P entails Q, we might say that if a person doesn't know that P entails Q, then that person can't competently deduce Q from P. Such an understanding of competent deduction, however, is implausible. Perhaps some ways of failing to know that P entails Q would constitute a failure of deductive competence, but this particularly devious way is not one of them. When people conspire to rob you of knowledge of an entailment, they sometimes will succeed at that, but that success doesn't rob you of the ability in question. To see this last point, just notice that any test that might be given for presence of the ability would still be one on which you perform flawlessly.

### Works Cited

- Audi, R. "Deductive Closure, Defeasibility and Scepticism: A Reply to Feldman." *Philosophical Quarterly* 45 (1995): 494–9.
- Bogdan, Radu J. "Cognition and Epistemic Closure." *American Philosophical Quarterly* 22 (1985): 55–63.
- Brueckner, Anthony. "Deductive Closure Principle." *Routledge Encyclopedia of Philosophy*. Ed. E. Craig. London: Routledge, 1998. <<http://www.rep.routledge.com/article/P011>>.
- . "Skepticism and Epistemic Closure." *Philosophical Topics* 13 (1985): 89–117.
- . "Strategies for Refuting Closure for Knowledge." *Analysis* 64 (2004): 333–5.
- Christensen, David. *Putting Logic in its Place: Formal Constraints on Rational Belief*. Oxford: Oxford University Press, 2004.
- Cohen, Stewart. "Basic Knowledge and the Problem of Easy Knowledge." *Philosophy and Phenomenological Research* 65 (2002): 309–29.
- . "How to be a Fallibilist." *Philosophical Perspectives* 2 (1988): 91–123.
- . "Two Kinds of Skeptical Argument." *PPR* 58 (1998): 143–59.
- DeRose, Keith. "Solving the Skeptical Problem." *Philosophical Review* 104 (1995): 1–52.
- Dretske, Fred. "The Case Against Closure." *Contemporary Debates in Epistemology*. Eds. M. Steup and E. Sosa. Malden, MA: Blackwell, 2005. 13–25.
- . "Conclusive Reasons." *Australasian Journal of Philosophy* 49 (1971): 1–22.
- . "Epistemic Operators." *Journal of Philosophy* 67 (1970): 1007–23.
- Feldman, Richard. "In Defense of Closure." *The Philosophical Quarterly* 45.181 (1995): 487–94.
- Goldman, Alvin. "Discrimination and Perceptual Knowledge." *Journal of Philosophy* 73 (1975): 771–91.
- Hales, Steven D. "Epistemic Closure Principles." *Southern Journal of Philosophy* 33 (1995): 185–201.
- Harman, Gilbert. *Change in View*. Cambridge, MA: MIT Press, 1989.
- Hawthorne, John. "The Case for Closure." *Contemporary Debates in Epistemology*. Eds. M. Steup and E. Sosa. Malden, MA: Blackwell, 2005. 26–42.
- . *Knowledge and Lotteries*. Oxford: Oxford University Press, 2004.
- Heller, Mark. "Relevant Alternatives and Closure." *Australasian Journal of Philosophy* 77 (1999): 196–208.
- Klein, Peter. "Closure Matters: Academic Skepticism and Easy Knowledge." *Philosophical Issues* 14 (2004): 165–84.
- . "Skepticism and Closure: Why the Evil Demon Argument Fails." *Philosophical Topics* 23 (1995): 213–36.
- Kvanvig, Jonathan L. "Contextualism, Contrastivism, Relevant Alternatives, and Closure." *Philosophical Studies* (forthcoming 2006).
- Kyburg, Henry. *Probability and the Logic of Rational Belief*. Dordrecht: Kluwer, 1961.
- . "Probability and Randomness." *Theoria* 29 (1963): 27–55.
- Lewis, David. *Counterfactuals*. Boston, MA: Blackwell, 1973.
- Luper, Steven. *Stanford Encyclopedia*. <<http://plato.stanford.edu/entries/closure-epistemic/>>.
- Maitzen, Steve. "The Knower Paradox and Epistemic Closure." *Synthese* 114 (1998): 337–54.
- Makinson, D. C. "The Paradox of the Preface." *Analysis* 25 (1965): 205–7.
- Nozick, Robert. *Philosophical Explanations*. Cambridge, MA: Harvard University Press, 1981.
- Pritchard, Duncan. "Closure and Context." *Australasian Journal of Philosophy* 78 (2000): 275–80.
- . "McKinsey Paradoxes, Radical Scepticism, and the Transmission of Knowledge across Known Entailments." *Synthese* 130 (2002): 279–302.

- . “Radical Scepticism, Epistemological Externalism and Closure.” *Theoria* 69 (2002): 129–61.
- . “The Structure of Sceptical Arguments.” *The Philosophical Quarterly* 55 (2005): 37–52.
- Ramsey, F. P. “Knowledge.” *FP Ramsey: Philosophical Papers*. Ed. D. H. Mellor. Cambridge: Cambridge University Press, 1990. 110–11.
- Shope, Robert K. *The Analysis of Knowing: A Decade of Research*. Princeton: Princeton University Press, 1983.
- Stalnaker, Robert. “The Problem of Deduction.” *Inquiry*. Cambridge, MA: MIT Press, 1984. 79–101.
- Stine, G. C. “Dretske on Knowing the Logical Consequences.” *Journal of Philosophy* 68 (1971): 296–9.
- . “Skepticism, Relevant Alternatives, and Deductive Closure.” *Philosophical Studies* 29 (1976): 249–61.
- Vogel, Jonathan. “Are there Counterexamples to the Closure Principle?” *Doubting*. Eds. M. D. Roth and G. Ross. Dordrecht: Kluwer Academic Publishers, 1990. 13–27.
- Warfield, Ted A. “When Epistemic Closure Does and Does Not Fail: A Lesson from the History of Epistemology.” *Analysis* 64 (2004): 35–41.
- and Marian David. “Knowledge-Closure and Skepticism.” *Epistemology: New Essays*. Ed. Quentin Smith. (forthcoming).
- Williamson, Timothy. *Knowledge and Its Limits*. Oxford: Oxford University Press, 2000.
- Wittgenstein, Ludwig. *On Certainty*. Eds. G. E. M. Anscombe and G. H. von Wright. Trans. Denis Paul and G. E. M. Anscombe. Oxford: Oxford University Press, 1969.