

PART I

KNOWLEDGE AND SKEPTICISM

Introduction

Matthias Steup

Chapter 1: Is Knowledge Closed under Known Entailment?

Sometimes our beliefs are true by chance. In such cases, they do not amount to knowledge, for if we know what we believe, it must not be by accident that what we believe is true. How can we capture what it takes to rule out such accidentality? According to Fred Dretske (1971), what rules out accidentality is the possession of conclusive reasons. Thus Dretske holds that knowing that p requires having a conclusive reason for p . He defines the concept of a conclusive reason counterfactually as follows:

Conclusive Reason

R is a conclusive reason for $p =_{df}$ If p were false R would be false (or: If p were not the case R would not be the case).

Counterfactuals are claims about what would (or would not) be the case if things were different than they in fact are. The standard method of assessing the truth value of a counterfactual is to see whether its consequent is true in the closest possible worlds in which its antecedent is true.

Suppose, prompted by a goat-in-the-field-like visual experience, E , Jim forms the belief

G : There is a goat in the field.

On Dretske's view, Jim knows G only if his reason for G , E , is conclusive. To assess whether E is a conclusive reason for G , we must determine the truth value of the following counterfactual:

C1: If there were no goat in the field, Jim would not have E.

We could imagine fancy circumstances in which the consequent of C1 is false. For example, a Cartesian demon might manipulate Jim's brain so as to let him have E even though there is no goat in the field. But we must consider the closest possible world in which G is false. That would be a world in which everything remains unaltered except for one thing: there is no goat in the field, but instead perhaps a cow, or no such animal at all. In such a world, Jim would not have E. Instead, he would have a visual experience of an empty field, or a field with a cow in it. When assessing C1 against a world like that, it turns out true. So E is a conclusive reason for G, which is to say Dretske's theory yields the result we ought to get: G is a proposition that Jim knows.

Thinking of knowledge as requiring conclusive reasons, as Dretske does, has an important consequence: the principle that knowledge is closed under known implication – *Closure*, for short – turns out to be false. Exactly how the Closure Principle is to be articulated is a matter of some complication. A rough and ready version of it goes as follows:

Closure

If S knows that p and knows that p implies q, then S is in a position to know that q.

There is broad agreement that deducing logical consequences of our knowledge can always expand our knowledge, i.e. can always result in knowledge of these consequences. There are, however, detractors. Dretske is one of them. For if knowledge requires conclusive reasons, it turns out that sometimes we are not in a position to know the logical consequences of what we know. Here is a well known example from Dretske (1970). Visiting the zoo, you see several zebras in an enclosure. Your zebra-like visual experience, E_z , is a conclusive reason for, and thus gives you knowledge of

Z: There are zebras in the enclosure.

What makes E_z conclusive is this: the closest possible world in which Z is false is one in which there are, say, tigers in the enclosure. In such a world, you would not have E_z . Now consider the following alternative, incompatible with Z:

M: The animals in the enclosure are cleverly disguised mules.

The question is: do you know that M is false, i.e. do you have a conclusive reason for $\neg M$? As long as your evidence for Z consists in E_z only, you do not. The relevant counterfactual is

C2: If M were true, I would not have E_z .

Since cleverly disguised mules look exactly like zebras, you *would* have zebra-like visual experiences if M were true. So C2 is false. Hence you do not know $\neg M$ even though you know that Z implies $\neg M$. Closure fails.

In his contribution to this volume, Dretske makes a case for embracing this result. Closure is bound to fail, he argues, because the ordinary things we know have heavyweight implications that are unknowable. For example, you can know that the following event took place at some time in the past: John ate all the cookies. This event implies the following *heavyweight propositions*:

There are physical objects.
There are other minds.
The past is real.

According to Dretske, there are no accredited ways of knowing such propositions. Perception can give you knowledge about external objects, but it cannot give you conclusive reasons for thinking that idealism is false. For if idealism were true, your perceptual reasons would remain the same. Your memory can give you knowledge of the past, but it can't give you conclusive reasons for thinking that the past is real. For if the world, complete with fossils and memories of the past, had come into existence five minutes ago, your memorial reasons would remain unaltered. In general terms, instruments can be reliable indicators of certain features of reality. But they cannot be indicators of their own reliability, no matter how reliable they may in fact be. It is difficult to see, then, how we could know the heavyweight implications of our ordinary knowledge. Hence we face a trilemma. It consists of the following three propositions:

- A: We have lots of knowledge of ordinary things.
- B: There is no way of knowing the heavyweight implications of our knowledge of ordinary things.
- C: Closure.

We know that the ordinary things we take ourselves to know imply heavyweight propositions of the kind mentioned above. Thus, if B and C are true, A must be false. But unless we are attracted to skepticism, certainly we do not want to give up A. If A and C are true, B must be false. According to Dretske, rejecting B amounts to verbal *hocus pocus*. To know that the heavyweight implications of our ordinary knowledge are true, we must know that the following *skeptical alternatives* are false:

- We are deceived by a Cartesian evil demon.
- Solipsism: I am the only conscious thing in existence.
- The world sprang into existence five minutes ago, complete with fossils and memories of the past.

Each of these alternatives is designed so as to ensure that our evidence remains unaltered. As a result, we cannot have conclusive reasons for their negations. Hence, on Dretske's view, we cannot know them to be false.

So Dretske considers the denial of neither A nor B an attractive option for resolving the trilemma, and hence views Closure as the sole remaining choice for rejection. Compared with A or B, denying C, Dretske suggests, carries the smallest costs. After

all, abandoning Closure does not mean that we can *never* use implications to expand our stock of knowledge. Frequently, when we know that p implies q , we *will* have conclusive reasons for both p and q . In cases like that, we are in a position to deduce q from p .

In his essay in defense of Closure, John Hawthorne begins by discussing how a plausible version of Closure is to be articulated, and then advances three main criticisms of Dretske's view. First, rejecting Closure turns out to be more costly than it may initially seem. For if Closure fails, so does the following, rather plausible principle (Distribution): if you know the conjunction of p and q and know that a conjunction implies each of its conjuncts, then you are in a position to know p (or, for that matter, q). The other two criticisms concern Dretske's distinction between ordinary propositions for which we can have conclusive reasons, and heavyweight propositions for which we can't have such reasons. According to Hawthorne's second objection, there are some heavyweight propositions for which it is possible to have conclusive reasons. For example, although you *cannot* have a conclusive reason for the proposition that you are not a brain in a vat (a BIV), you *can* have a conclusive reason for the conjunction *I have a headache and am not a BIV*. So, on Dretske's view, you cannot know that you are not a BIV, but you can know that you have a headache and are not a BIV. Third, Hawthorne argues that there are cases in which a subject fails to have a conclusive reason for a proposition that we would classify as easily knowable. His example involves various quantities of consumed salmon. It is easy to have a conclusive reason for believing that you consumed less than one pound of salmon. But if consuming a very large quantity of salmon (say 14 pounds) causes you to hallucinate that you ate very little salmon, you do not have a conclusive reason for believing that you ate less than 14 pounds of salmon. So if you suffer from this affliction, you can know that you ate less than one pound of salmon, but not that you ate less than 14 pounds. In sum, then, Hawthorne argues that conceiving of knowledge in terms of conclusive reasons bears significant costs.

Dretske's reply is simple. A convincing defense of Closure must show that there is an alternative account of knowledge that (a) succeeds in fending off skepticism and (b) does so at smaller costs than Dretske's conclusive reasons account. Hawthorne, Dretske says, has not done that. Without endorsing it, Hawthorne reviews contextualism as an alternative way of defending our ordinary knowledge against skepticism. Dretske suggests that, when comparing costs, contextualist grass is in worse shape than the grass in his yard.

Chapter 2: Is Knowledge Contextual?

Unlike Dretske, contextualists hold that Closure should be preserved. Yet they agree with Dretske that the falsehood of a skeptical alternative such as "I'm a BIV" cannot be known – at least according to the standards of knowledge in place when we worry about skeptical alternatives. But if Closure is true *and* it is true that one cannot know that one is not a BIV, then we are confronted with what appears to be an inconsistent triad of very plausible propositions:

- 1 I know that I have hands.
- 2 I do not know that I am not a BIV.
- 3 I know that I have hands \rightarrow I know that I am not a BIV.

It is hard to see how I could fail to know that I have hands. It is also hard to see how I could know that I am not a BIV, since if I were a BIV, my evidence would not be relevantly different from what it is now. Contextualists hold, therefore, that 1 and 2 enjoy a high degree of plausibility. Now, we could secure 1 and 2 by rejecting 3. But contextualists wish to preserve Closure. Since I know that my having hands implies my not being a BIV, Closure tells us that 3 is true. But if 3 is true, either 1 or 2 must be false. It would appear, then, that if we accept Closure, we must choose between 1 and 2. Contextualists stress that simply denying either 1 or 2 is not going to be a satisfying solution to the puzzle. What we need is a solution that preserves the intuitions in light of which 1 and 2 seem both true. According to contextualism, there is such a solution. It consists of an explanation of how propositions 1–3 can all be true (Cohen, 1988, 1999; DeRose, 1995).

Sometimes two assertions seem to be inconsistent when in fact they are not. Here is an elementary example to illustrate this phenomenon, involving the indexical “I”:

- 4 “I have a headache” . . . uttered by Jack.
- 5 “I do not have a headache” . . . uttered by Jill.

The conjunction of 4 and 5 seems to be of the form “p and \neg p,” and thus inconsistent. However, a sentence such as “I have a headache” has different truth conditions when it is uttered by different speakers. If Jack has a headache but Jill does not, “I have a headache” is true when uttered by Jack and false when uttered by Jill. Consequently, if Jack has a headache and Jill does not, 4 and 5 are both true.

The word “knowledge,” contextualists say, functions just like an indexical. The truth conditions of a sentence such as “I am F” depends on who asserts it. According to contextualists, the truth conditions of sentences such as “S knows that p” and “S does not know that p” depend on the context of the subjects who assert them. Let us refer to sentences that attribute knowledge (or ignorance) to a subject as *K-assertions*. Contextualists hold that whether a K-assertion is true depends on which standard of knowledge is in place in the context in which the assertion is made. In some contexts, the standards of knowledge are low; in others, they are high. In ordinary situations of daily life, the standards of knowledge are less demanding than they are in philosophical contexts. When philosophers get together and worry about skeptical alternatives, the standards rise to the maximum: the assertion that S knows that p will be true only if S can eliminate *any* alternative that is incompatible with p. In the various non-philosophical contexts of ordinary life, the standards are lower. In low-standard contexts, what is required for the assertion that S knows p to be true is merely that S can eliminate alternatives to p that are *relevant* in those contexts.

Let us see how contextualism works when we apply it to an example:

- 6 Alan Greenspan knows he has hands.
- 7 Alan Greenspan does not know he has hands.

On the face of it, 6 and 7 seem inconsistent. Contextualists hold that they need not be. Assume Jack and his friend Joe are in a bar debating, while having a beer together, what Alan Greenspan knows and what he doesn't know. Jack claims that Greenspan knows that he has hands, but doesn't know when the Dow will reach 12,000. Since neither Jack nor Joe worries about being a BIV, or any other skeptical alternative, the standards of knowledge remain low. In that context, Jack's assertion of 6 is true.

Assume further that Jill is attending a philosophy seminar on skepticism. In this context, the skeptical alternative of being a BIV becomes salient. As a result, the standards of knowledge are high. Just like Jack and Joe, Jill happens to be thinking about what Alan Greenspan knows. Noting that Greenspan cannot eliminate the possibility that he is a BIV, Jill claims that Greenspan does not know that he has hands. In that context, her assertion of 7 is true. Since Jack's assertion of 6 and Jill's assertion of 7 occur in different contexts, each of them involving a different standard of knowledge, both assertions turn out to be true.

So, according to contextualism, whether an attribution of knowledge is true depends on the standards in place in the attributor's context. To determine the truth value of a K-assertion, we must determine whether: (a) the standards in place in the attributor's context are high or low; and (b) the subject to whom knowledge is attributed meets these standards.

Let us return to the puzzle that arises from accepting Closure. If Closure is true, it would seem that either

1 I know that I have hands.

or

2 I do not know that I am not a BIV.

is false. According to contextualism, we do not have to choose between 1 and 2. In low standard contexts, we can truly assert 1 and the negation of 2; in high standard contexts, we can truly assert 2 and the negation of 1.¹ So according to the contextualist solution to the puzzle, we can preserve both our skeptical and non-skeptical intuitions. This, according to Cohen, is one of the virtues of contextualism.

What theory does contextualism compete with? In his critique of contextualism, Earl Conee argues that its plausibility must be gauged against that of *high standards invariantism*. According to this view, the truth conditions of K-assertions do not vary from context to context. Instead, there is one, and only one, standard: one that is high, but not excessively high. Call this the *unique standard*. In ordinary life, we frequently attribute knowledge to subjects who do not have knowledge according to the unique standard. When this happens, we are simply indulging ourselves in loose talk. On the other hand, when, for some p we would normally claim to know, we are confronted with the question of whether we *really* and *truly* know p, we tend to apply the unique standard and, if convinced we don't meet that standard, retract the knowledge claim. We will not, however, retract K-assertions when it comes to the best cases of perception, memory and triple checked calculations. Such cases satisfy the unique

standard, and thus are genuine instances of knowledge. Skepticism denies that they are instances of knowledge, and must therefore be rejected.

Earl Conee defends high standards invariantism on the following grounds. First, it preserves conflict where genuine conflict exists. Sceptics deny our ordinary claims to knowledge. Thus we judge that propositions 1–3 form an inconsistent set. If contextualism is correct, this judgment, the truth of which should be preserved, turns out to be mistaken. Second, contextualism concedes too much to skepticism. Like other philosophical arguments, those for skepticism can be critically assessed and might be judged as wanting. They do not deserve the deference that contextualism affords them. Third, if contextualism were true, we would be faced with another paradox. If the truth conditions of K-assertions really were context-sensitive, this should be a semantic fact transparent to competent speakers. Why, then, do competent speakers fail to recognize the context sensitivity of “know,” and why does contextualism elude broad acceptance within the philosophical community?

Stewart Cohen defends contextualism on the ground that it offers us a better resolution of the paradox than high standard invariantism does. In response to Conee, he makes the following points. First, when responding to the paradox, we will have to acknowledge error somewhere. It is correct that, if contextualism is true, we are mistaken in thinking that propositions 1–3 are inconsistent. But if invariantism is correct, we must judge one of these propositions to be false, although each of them is highly plausible. Second, Cohen denies that contextualism concedes too much to skepticism. Sceptical arguments are cogent arguments. They do not suffer from obvious defects. It is a virtue of contextualism that it captures what is plausible about skepticism and does so in a way that preserves our ordinary claims to knowledge. Third, Cohen argues that “know” is not the only word the context-sensitivity of which escapes competent speakers. According to flatness skepticism, nothing is really flat because no surface is entirely free of bumps. When confronted with flatness skepticism, competent speakers respond not by noticing the context-sensitivity of “flat,” but by becoming convinced that, indeed, nothing is flat. Finally, Cohen suggests that the explanation of why the philosophical community did not rush to endorse contextualism is that it is a good news–bad news theory. The good news is that, in ordinary contexts, we can truly attribute to ourselves the kind of knowledge we ordinarily think we have. The bad news is that, in philosophical contexts, we cannot truly assert that we know we have hands. Most philosophers don’t like to acknowledge this, so they reject contextualism.

Chapter 3: Can Skepticism Be Refuted?

I know that my having hands implies my not being a BIV. Suppose further that we accept Closure. It then follows that I either know that I’m not a BIV, or do not know that I have hands. So if we accept Closure, we face the following skeptical argument:

The BIV Argument

- 1 If I know that I have hands, then I know (or am at least in a position to know) that I am not a BIV.

- 2 I do not know that I'm not a BIV.
 Therefore:
 3 I do not know that I have hands.

The first premise is the result of accepting Closure. So if you accept Closure, and you think you know you have hands, you must find a good objection to 2. Dretske, as you will recall, thinks that any such objection will amount to verbal *hocus pocus*. According to Jonathan Vogel, however, there are good grounds for rejecting 2.

Suppose you don't know how to tell a goldfinch from a canary. Vogel would say that, if so, your choice between

G: This bird is a goldfinch

and

C: This bird is a canary

is evidentially underdetermined. Your evidence is not sufficient for making a justified choice between C and G. The same, the skeptic would argue, applies to the two propositions:

- H: I have hands.
 B: I'm a BIV.

The BIV hypothesis is designed so as to make your evidence compatible with your being a BIV. No part of your evidence, the skeptics would argue, gives you a clue as to whether or not you are a BIV. Hence your evidence underdetermines your choice between H and B. That is why you cannot know that you are not a BIV. Vogel suggests that the choice between H and B is in fact not underdetermined. You do have reasons that you could bring to bear against B.

To reject the underdetermination claim, Vogel appeals to the principle of *inference to the best explanation*: if hypothesis A explains the relevant data better than hypothesis B, then we have a reason to prefer A to B. When discussing skepticism, the relevant data are our ordinary perceptual experiences. Vogel suggests that undeveloped skeptical hypotheses – hypotheses that do not explain in detail how our experiences come about – are poor explanations, clearly inferior to the physiological explanation we can employ when we assume that the world is pretty much what we take it to be (the *real world hypothesis*, in Vogel's terminology). However, skeptical hypotheses need not suffer from a paucity of detail and a lack of sophistication. What Vogel calls the *isomorphic skeptical hypothesis* is designed to match the cause and effect relationships of the real world hypothesis. Its explanatory power seems therefore to be no less than that of the real world hypothesis. On which grounds, then, could it be rejected?

According to Vogel, the isomorphic skeptical hypothesis suffers from the following defect: though it offers us a more complex explanation of the relevant data than the real world hypothesis, it does not deepen our understanding of these data. Its

increased complexity is not offset by a commensurate increase in explanatory scope and depth. But why is the isomorphic hypothesis more complex? It is more complex, according to Vogel, because it must replace genuine shapes with pseudo shapes, and genuine locations with pseudo locations. That is why we may prefer the real world hypothesis after all, and thus reject the second premise of the BIV Argument.

Unlike Vogel, Richard Fumerton does not attempt a refutation of skepticism. Instead, he examines the nature of the challenge before us when we attempt to avoid skepticism within the traditional framework of foundationalist, internalist philosophy. If we accept this framework, then, according to Fumerton, we are committed to a view that he calls *inferential internalism*. It can be stated as follows:

- (i) We are acquainted with the contents of our thoughts, and thus have direct knowledge of them.
- (ii) We are not acquainted with physical objects. Therefore, if we have knowledge of physical objects, it must be inferential: inferred from what we know directly on the basis of acquaintance.
- (iii) Call a body of evidence obtained through acquaintance “E.” Call a typical proposition about an ordinary object “O.” If we can acquire knowledge of O by inferring it from E, E must make O probable, and we must *know* that E makes O probable.

Many philosophers would reject (iii). To such philosophers, Fumerton would reply as follows. Suppose a subject, S, makes a claim, C, on the basis of a body of evidence, E. Suppose further that you are given the information that S has no clue as to whether or not E makes C probable. Wouldn't you then judge that S is *not* justified in claiming C to be true, even if you are yourself convinced that E does in fact make C probable? If you would thus judge, this shows that you actually accept condition (iii).

One problem with condition (iii) is that philosophers who accept it face a regress problem. How could we ever know a proposition to the effect that E makes O probable? Suppose we acquire knowledge of

P1: E makes O probable

by inferring it from a proposition P2. In that case, (iii) requires us to know

P3: P2 makes P1 probable.

If our knowledge of P3 were inferential as well, we would have to infer P3 from a proposition, P4, and we would have to know that P4 makes P3 probable, and so forth. So if knowledge of probability were inferential, we would end up in an infinite regress. This would make it rather doubtful that such knowledge can be had. Thus the inferential internalist must, as Fumerton stresses, employ a Keynesian conception of probability, according to which we can know directly that E makes O probable. According to Keynes, knowledge of probability can be understood in analogy to knowledge of entailment. For example, we can know directly that “p” entails “p or q.” Likewise,

according to Keynes, if E does indeed make O probable, we can know directly that this is so.

According to the inferential internalist, I cannot know *directly* that I am not a BIV. Thus it is not easy for the inferential internalist to avoid skepticism. If inferential internalism is correct, and if I am to know that I am not a BIV, there must be a body of evidence that makes my not being a BIV probable for me, and it must be possible for me to know directly that it makes this probable for me. That, according to Fumerton, is the challenge we face when we wish to refute skepticism.

Vogel's suggestion is that the needed body of evidence consists in the fact that the BIV hypothesis is not as good an explanation of the relevant data as the real world hypothesis. But does it? Suppose that, all other things being equal, the simpler of two explanations is more likely to be true than the more complex one. Obviously, antiskepticism wins only if the real world hypothesis is indeed simpler than any competing skeptical hypothesis. Fumerton points out, however, that Berkeley claimed his own view of immaterialism – according to which God causes our experiences – was simpler than the competing view that our experiences are caused by physical objects. Moreover, why should we believe that the simpler of two explanations is really more likely to be true? Fumerton argues that the skeptic has every right to question the probability principles to which the non-skeptic appeals. What reasons, then, can be adduced for thinking that inference to the best explanation is a guide to truth? Vogel responds that this does indeed raise a difficult problem for the kind of antiskepticism he advocates, but registers the opinion that inference to the best explanation is in fact no more problematic than any other principle of probability.

Chapter 4: Is There *a Priori* Knowledge?

There is a large set of propositions that, according to philosophical tradition, are knowable *a priori*: knowable independently of experience. What is meant by saying that such propositions are knowable in this way is that the justification we have for believing them does not arise from experience. Thus the issue over *a priori* knowledge really concerns a distinction between two kinds of *justification*: empirical and *a priori*, the former resulting from the use of our perceptual and introspective faculties, the latter from exercising what we might call the faculty of rational insight. Typical examples of the kind of propositions claimed to be knowable through rational insight are the following:

The sum of two and two is four.

Two is an even number.

Whatever is square is rectangular.

Whatever is red is colored.

Necessarily, if “p or q” is true and “q” is false, then “p” is true.

Call such propositions PAPS: putatively *a priori* propositions. It is generally agreed that PAPS are knowable. There is skepticism about knowledge of the external world, other minds, and the past. Skepticism about PAPS, however, is rarely pursued. Indeed,

considering that knowledge of PAPS includes knowledge of the laws of logic, and more specifically, knowledge of an argument's validity, it is hard to see how a skeptical argument for anything could get off the ground without the prior assumption that knowledge of PAPS is indeed possible. So the knowability of PAPS is not at issue. However, there are philosophers – advocates of radical naturalism – who question that PAPS are really known *a priori*. They can be viewed as skeptics about *a priori* knowledge, not on account of denying that we know any PAPS, but because they hold that there is no such faculty as rational insight, and consequently no such thing as *a priori* knowledge. According to such philosophers, all knowledge, including that of PAPS, is empirical.

In the two essays on the issue of *a priori* knowledge, Laurence Bonjour represents the traditional point of view, and Michael Devitt makes the case for radical naturalism. For defenders of *a priori* knowledge, the chief task is to explain how *a priori* justification comes about. According to Bonjour, justification for an *a priori* proposition, *p*, results from grasping that *p* is necessarily true. We recognize *p* as a necessary truth, and thus are supplied with a reason for accepting *p*. For opponents of *a priori* knowledge, the task is twofold. First, they must give an account of how empirical knowledge of PAPS is possible. Second, they must explain why PAPS cannot be known through rational insight. Devitt attempts to dispose of the first task by appeal to the thesis of holism, proposed by Duhem and Quine: beliefs face the tribunal of experience not individually but only as a system. PAPS are part of our belief system. Since any belief within such a system may be revised in light of experience, it follows that there is empirical justification for PAPS.

However, from the thesis of holism, it does not follow that, in addition to holistic empirical justification, there is not also *a priori* justification for PAPS in the form of rational insight. Thus Devitt offers a second, main argument. *A priori* justification in the form of rational insight is, he claims, utterly obscure. Although we lack a well developed theory of empirical justification, it is nevertheless well understood what such justification at bottom amounts to. But as far as rational insight is concerned, we face an unexplained mystery. Not only do we not have a well developed account of it, we do not even understand how we *could* acquire knowledge of external, worldly facts without the help of experience. In short, there is a significant discrepancy between empirical justification and *a priori* justification. The latter is deeply mysterious, the former not at all.

Bonjour denies such discrepancy. Empirical justification, at least when it comes to indirect empirical knowledge, is no less problematic than *a priori* justification. In response to Devitt's appeal to holism, Bonjour claims that Devitt begs the question. If the premise of holism is to yield the conclusion that PAPS are justifiable empirically, we must assume the additional premise that a conflict between a PAP and a set of experiences must always be decided in favor of the latter. But this is just what is at issue, and what a traditional philosopher would deny. Moreover, Bonjour argues that even if we grant the premise of holism, we still need *a priori* knowledge. According to a Quinean epistemology, there are certain features of belief systems that count in favor of their truth; for example, simplicity, fecundity, explanatory strength. How can we justify the claim that such features are really indicators of truth? If we cannot do this *a priori*, we must infer it within our holistic belief system. But any such

procedure would be viciously circular. It follows that the holistic naturalist cannot explain why holistic justification supplies us with a reason for thinking that the holistic belief system is true. Skepticism is the inevitable outcome.

BonJour, then, denies that holism can account for our knowledge of PAPS. In response, Devitt argues for three claims:

- (i) For holistic evidence to justify our beliefs, it is not necessary for us to know that the items that constitute holistic evidence are indicators of truth.
- (ii) The holistic naturalist is committed not to premise circularity, but to rule circularity only, which, unlike the former, is benign.
- (iii) BonJour's rationalist defense of apriority requires rule-circularity no less than holistic naturalism, and thus is committed to viewing such circularity as benign.

BonJour rejects each of these claims. First, he insists that an evidential item, E, can justify our acceptance of p only if we know that E is reason for thinking that p is true. Second, he denies that rule circularity is benign. Finally, he argues that *a priori* justification through rational insight does not raise circularity problems because it is immediate, non-propositional, and thus atomistic.

Note

- 1 It is, however, not entirely clear whether it is possible truly to deny 2 in a low standard context. Let "B" stand for the proposition that I am a BIV. Can I attribute to myself knowledge of $\neg B$ while remaining in a low standard context? Perhaps merely thinking about the envatment possibility makes the skeptical alternative of being a BIV salient, thus creating a high standard context. If this is so, it remains doubtful that one can ever *truly* attribute to oneself (or to anyone else, for that matter) knowledge of $\neg B$.

References

- Cohen, S. (1988) How to be a fallibilist. *Philosophical Perspectives*, 2, 581–605.
 Cohen, S. (1999) Contextualism, skepticism, and the structure of reasons. *Philosophical Perspectives*, 13, 57–89.
 Dretske, F. (1970) Epistemic operators. *Journal of Philosophy*, 67, 1007–23.
 Dretske, F. (1971) Conclusive reasons. *Australasian Journal of Philosophy*, 49, 1–22.
 De Rose, K. (1995) Solving the skeptical problem. *Philosophical Review*, 104, 1–52.

CHAPTER
O N E

Is Knowledge Closed under Known Entailment?

The Case against Closure

Fred Dretske

Closure is the epistemological principle that if S knows that P is true and knows that P implies Q, then, evidentially speaking, this is enough for S to know that Q is true. Nothing more is needed. If S believes Q on this secure basis – on the basis of two things he knows to be true – then S knows that Q is true. One knows everything that one knows to be implied by what one knows.

It is important to distinguish closure from *modus ponens*, a principle of logic that it superficially resembles. Closure is stronger. *Modus ponens* says that if P is true, and if P implies Q, then Q must be true. Closure tells us that when S knows that P is true and also knows that P implies Q, then not only must Q be true (*modus ponens* gets you this much), S must *know* it is true.¹

Why must S know Q to be true? Why does one have to know everything one knows to be implied by what one knows? One doesn't, after all, have to regret everything one knows to be implied by what one regrets. Tom regrets drinking three martinis last night, but he doesn't regret what he knows to be implied by this – drinking *something* last night. Nor does he regret the reality of the past even though he knows that drinking three martinis last night implies that the past is real. If

(R) S regrets P
S knows that P implies Q
 Therefore, S regrets Q

is not a valid argument, why should

(K) S knows P
S knows P implies Q
 Therefore, S knows Q

be accepted as valid?

This question defines the issue between John Hawthorne and me. I will be taking the negative side of the debate: knowledge is *not* closed under known logical implication. (K) is *not* a valid argument. There are things we know to be implied by what we know that we do not know to be true. Just as a great many conditions I do not regret have to exist (and I know they have to exist) in order for the condition I regret to exist, a great many conditions I do not know to exist have to exist (and I know they have to exist) in order for the conditions I know to exist.

1 Transmissibility

Let me begin this examination by talking about a problem that is closely related to closure: transmission of evidential warrant. I shall come back to closure later (in sections 2 and 3).

Our ways of discovering P are not necessarily ways of discovering what we know to be implied by P. From the fact that you know that P implies Q, it does not follow that you can see (smell, feel, etc.) that Q just because you can see (smell, feel, etc.) that P. Despite knowing that cookies are objective (mind-independent) objects, I can see (roughly: tell by looking) that there are cookies in the jar without being able to see, without being able to tell by looking, that there are mind-independent objects. A claim to have found out, by looking, that there are cookies in the jar is not a claim to have found out, by looking, that there is a material world. Maybe one has to know there are physical objects in order to see that there are cookies in the jar (we will come back to that), but one surely isn't claiming to see that there are physical objects in claiming to see there are cookies in the jar. After all, hallucinatory cookies "in" hallucinatory jars can look exactly like real cookies in real cookie jars. So one cannot, not by vision alone, distinguish real cookies from mental figments. One cannot see that the world really is the way it visually appears to be. A way of knowing there are cookies in a jar – visual perception – is *not* a way of knowing what one knows to be implied by this – that visual appearances are not misleading.

The claim to have found out, by visual means, that there is still wine left in the bottle ("Just look; you can see that there is") is not a claim to have found out, at least not by visual means, what you know to be implied by this – that it, the liquid in the bottle, is not merely colored water. You know it is wine because, let us suppose, you tasted it a few minutes ago. That you learned it is wine, and not merely colored water, by tasting does not prevent you from now (minutes later) seeing that there is still some wine left in the bottle. Seeing that P does not mean you can see that Q just because you know P implies Q.

In *Seeing and Knowing* (1969) I tried to describe this phenomenon by introducing a technical term: *protoknowledge*. Protoknowledge was a word I made up to describe things that had to be true for what you perceived to be true but which (even if you knew they had to be true) you couldn't perceive to be true. That there are material objects has to be true for there to be cookies in the jar (and, hence, for you to see that there are cookies in the jar), but it isn't something you see to be so. It is a piece of (what I was calling) protoknowledge. When you see, just by looking, that there is wine in the bottle, the fact that it is not colored water is also protoknowledge. It has

to be true for there to be wine in the bottle (what you perceive to be so), but it is not a fact that (normally) you can see to be so. In describing how one knows there is still wine left in the bottle one is not (not normally) describing how one knows it is not merely colored water.

Now, thirty years later, as a result of work by Davies (1998, 2000) and Wright (2003), there is, I think, a more revealing way of making this point. The idea is that some reasons for believing P do not *transmit* to things, Q, known to be implied by P. In normal circumstances – a wine tasting party, say – one’s reasons for thinking there is wine left in the bottle do not transmit to, they are not reasons for believing, what you know to be implied by this – that the liquid in the bottle is not merely colored water. Having already tasted it, you may know that it is wine and not just colored water, but the point is that your reasons for believing the one (visual) are not reasons for believing the other (gustatory). Colored water in the bottle would look exactly the same as the wine. The reasons you have for believing what you say you perceive (there is wine left in the bottle) are not *transmitted* to this known consequence (that it is not merely colored water) of what you perceive.

The non-transmissibility (to many of the known consequences) of most of our reasons for believing P is an absolutely pervasive phenomenon. I will, in this section, spend a little more time harping about it since I think it critical for evaluating the plausibility of closure. Non-transmissibility does not itself imply the failure of closure since, as our wine example illustrates, even when S’s reasons for believing P do not transmit to a known consequence, Q, it may be that S must still know Q (perhaps on the basis of other reasons) in order to know P. Even though S cannot see that it is wine, not just colored water, it may turn out (this is what closure tells us) that S must know (perhaps by earlier tasting) that it is not colored water if he is now to see (hence, know) that there is wine left in the bottle. Nonetheless, once one appreciates the wholesale failure of evidential transmission, the failure of closure is, if not mandatory, easier to swallow. Or so I will argue.

As we have already seen, perception, our chief (some would say our only) route to knowledge of the world around us, does not transmit its evidential backing to *all* the known consequences of what is perceived. We can see (hear, smell, feel) that P, but some of the Qs that (we know) P implies are just too remote, too distant, to inherit the positive warrant the sensory evidence confers upon P. When Jimmy peeks into the cookie jar and, to his delight, sees that there are cookies there, his visual experience of the cookies, the evidential basis for his knowledge that there are cookies there, is not evidence, not a reason to believe, that there is a physical reality independent of Jimmy’s mind. Jimmy’s experience of the cookies may be good reason to believe there are cookies in the jar, but it is not a good reason to believe that idealism is false. And it is not a good reason to believe that idealism is false even if Jimmy understands that cookies are mind-independent objects and that, therefore, what he sees to be the case (that there are cookies in the jar) implies that idealism is false. Looking in the cookie jar may be a way of finding out whether there are any cookies there, but it isn’t – no more than kicking rocks – a way of refuting Bishop Berkeley.

So perceptual reasons – the sense experiences on which we base everyday perceptual judgments – do not transmit their evidential force to all the known consequences of the judgments they warrant.

- (T) R is a reason for S to believe P
S knows that P implies Q
 R is a reason for S to believe Q

is not a valid argument. When we perceive that P, there are, so to speak, *heavyweight* implications of P that cannot be perceived to be so. There is transmission to *lightweight* implications, of course. If I can see that there are cookies in the jar, I can certainly see that the jar isn't empty, that there is something in the jar. For perception, though, there are *always* heavyweight implications, known implications to what one perceives (P) that one's perceptual reasons for P are powerless to reach. If there is any doubt about this, simply imagine Q to be a condition S knows to be incompatible with P but which, because of (perhaps extraordinary) circumstances, has the same sensory effects on S as P. Though incompatible with P, Q (as so specified) will look (feel, sound) exactly the same to S as P. Q is, as it were, a perceptual twin of P. With such a Q, S will not be able to perceive $\sim Q$ though he perceives P and knows that P implies $\sim Q$.

Skeptics have used this formula to manufacture heavyweight implications for all our ordinary perceptual claims. No matter what you purport to know by perception, perception will not be the way you know that you are not being deceived by a Cartesian demon. It will not be the way you know you are not a disembodied brain in a vat being caused (by suitably placed electrodes) to experience the things an ordinary veridical perceiver experiences. One cannot see (hear, smell, feel, etc.) that one is not in this unfortunate position even though one knows one cannot be in this position and still be seeing (hearing, etc.) the things one believes oneself to be seeing (hearing, etc.). You can't see that you are not dreaming.

2 Closure

What does this have to do with closure? When you know that P implies Q, closure tells us you have to know Q to know P. It doesn't tell us that the reasons that promote your belief that P into knowledge that P must themselves promote your belief (assuming you have it) that Q into knowledge that Q. Maybe the reasons (see section 1) for P don't transmit to Q. Very well. Then closure tells us that if you know P you have to have *other* reasons (or, if reasons are not always required, whatever else it takes) to promote your belief that Q into knowledge that Q. But – or so closure says – you have to have *something* to make your belief that Q into knowledge in order to know P. If you don't, then – too bad – you don't know that P. Nothing I have said so far challenges this claim.

Is this true? Is it true that someone (who knows that cookies are material objects) cannot see that there are cookies in the jar unless he or she knows that there is a physical, mind-independent, world? That they are not being misled by some clever deception? In Dretske (1970) I asked the same question about my ability to see that an animal is a zebra without knowing it was not a cleverly painted mule. It seemed to me then, as it still seems to me today, that if this is true, if knowledge is closed under known implication, if, in order to see (hence, know) that there are cookies in

the jar, wine in the bottle, and a zebra in the pen, I have to know that I am not being fooled by a clever deception, that the “appearances” (the facts on which my judgments are based) are not misleading, then skepticism is true. I never see (to be so), and hence do not know (to be so), what I think and say I see (to be so).² This is not to deny that we know – and, perhaps, necessarily know – *most* of the implications of what we know. But not *all* of them. There are some things – *heavyweight* implications – we needn’t know even though we know our knowledge depends on their truth.

Though I haven’t performed a scientific study, my impression is that most philosophers (who bother thinking about it) believe that this is preposterous. Feldman (1999) thinks that abandoning closure is “one of the least plausible ideas to gain currency in epistemology in recent years.” DeRose (1995) finds it “intuitively bizarre” or “abominable.” Fumerton (1987) thinks the failure of closure is a “devastating objection” and Bonjour (1987) a *reductio ad absurdum* to any theory that implies or embraces it. Most philosophers, of course, are eager to reject skepticism. Rejecting closure as a way around skepticism, though, is quite another matter. As these reactions indicate, many philosophers are unwilling to do it. Skepticism is bad, yes, but for many philosophers, not *that* bad. Not bad *enough* to justify rejecting closure. That would be like rejecting logic because it gave us conclusions we didn’t like.

Some of these reactions are, I think, a bit overdone. To deny closure is not to say that you can *never* know (find out, discover, learn) that Q is true by inferring it from a P that you know to be true. It is merely to deny that this can be done for *any* Q. Nor is denial of closure a way of embracing (without further explanation) DeRose’s (1995) abominable conjunctions. Statements like “I know I have hands, but I do not know that I am not a handless brain in a vat” (a statement that denial of closure accepts as being possibly true) are clearly ridiculous. We can all agree about that. The question, however, is not whether anyone would ever say it (or, if they did, whether they would be greeted by stares of incomprehension) but whether it might be true if they said it. “The refrigerator is empty, but has lots of things in it” is also an abominable conjunction. It might, nonetheless, be true. There may be no food or other items normally stored in refrigerators inside (thus making the refrigerator empty in the normal way of understanding what isn’t in empty refrigerators), but it may, nonetheless, be filled with lots of gas molecules (perfectly respectable things). The abomination in saying a refrigerator is empty but has lots of things in it comes not from any logical defect (there are lots of empty things that have gas molecules in them – pockets, classrooms, warehouses, etc.), but from a violation of normal expectations. In describing an object as a refrigerator (and not, say, a metal container) one is led to expect that the things that are in (or, in the case of its being empty, not in) it are the sorts of perishable items normally stored or preserved in refrigerators. To then include (second conjunct) gas molecules as things in refrigerators is to flout this entirely reasonable expectation about what sorts of things are to be counted as things for purposes of describing the contents of refrigerators. Why isn’t it ridiculous, for exactly the same reason, to say one knows one has hands but doesn’t know one isn’t a handless brain in a vat? The second conjunct introduces possibilities normally assumed to be irrelevant (not counted as possibilities) by someone who asserts the first conjunct.

This is only to say (Grice 1967) that there are logical abominations (self-contradictory) and conversational abominations (perfectly consistent, and therefore possibly true statements, that violate conventional expectations). To say that S knows that there are cookies in the jar but doesn't know he isn't hallucinating them is certainly to say something absurd, but why suppose its absurdity is such (i.e. logical) as to render the statement incapable of being true? To demonstrate logical incoherence would require a theory about what it takes (and doesn't take) to know something, and we are back to where we started: assessing the status of closure.

So I don't think we (e.g. John Hawthorne and I) can just match brute intuitions on closure. Yes, closure sounds like an eminently plausible principle. Everything else being equal, then, we ought to keep it. But everything else *isn't* equal. It is also plausible (some would say it was entirely obvious) that we know things about our material surroundings – that (when we see them for ourselves) there are cookies in a jar, zebras in the zoo, people in the room, and cars on the street. And it isn't at all clear that we can have such knowledge while, at the same time, retaining closure. Something has to give. This reasoning won't impress a skeptic, of course, but, if a case could be made for the claim that a rejection of closure was not just *a* way to avoid skepticism (most philosophers would agree with this) but the *only* way to avoid skepticism, it should carry weight with philosophers who find skepticism as “bizarre” or “abominable” as the denial of closure. This will be my strategy in the remainder of this paper. The only way to preserve knowledge of homely truths, the truths everyone takes themselves to know, is, I will argue, to abandon closure.

Philosophers, though, will insist that there are other, less costly, ways of dealing with skepticism, ways that do not require the rejection of closure. One of the popular maneuvers is to make knowledge attributions sensitive to justificational context. In ordinary contexts in which knowledge claims are advanced, one knows (such things as) that there are cookies in the jar and (at the zoo) zebras in the pen. In such ordinary contexts one normally has all the evidence needed to know. One sees the cookies and the zebras and, given the absence of any special reason to doubt, this (for reasonably experienced adults) is good enough to know. So skepticism is false. Do these reasonably experienced adults also know, then, that they are not being deceived by some kind of fake? That the sensory evidence on which they base their judgments is not misleading? That there really is a material world? That they are not a brain in a vat? Well, yes (so closure is preserved), but – and here is the kicker – only as long as they don't seriously consider these heavyweight implications or say that they know them to be true. For once they think about them or say they know them to be true, the context in which knowledge claims are evaluated changes and knowledge evaporates. In this altered context (no longer an ordinary context) one doesn't know that one is not being deceived because new alternatives (that one *is* being deceived), possibilities one cannot evidentially eliminate, have been introduced. Therefore (closure forces one to say) one no longer knows that there are cookies in the jar and zebras in the pen. One gets to know about cookies and zebras only as long as one doesn't think about or claim to know what, according to closure, one has to know in order to know such mundane things. According to this way of dealing with skepticism, philosophers who spend time worrying about heavyweight implications (How do I know I'm not dreaming? How do we know there is a material world?) are the most

ignorant people in the world. Not only don't they know these heavyweight implications (maybe no one does), they don't (like everyone else) know the things (that there are cookies in the jar, zebras in the pen) that imply them. This, of course, includes almost everything they thought they knew.

This result, it seems to me, is pretty bizarre – *more* bizarre, in fact, than abandoning closure. It is a way of preserving closure for the heavyweight implications while abandoning its usefulness in acquiring knowledge of them. One knows (or is evidentially positioned to know) heavyweight implications only so long as one doesn't think about them or say that one knows them to be true. I think we can do better than this. By suitable restriction of closure (to lightweight and middleweight implications) we can have most of our cake (the only part worth eating) and eat it too.

Before launching this argument, though, let me mention a theoretical consideration that deserves attention in this regard. It is not, in the end, decisive on the issue of closure, but it is certainly relevant and with some readers it may carry weight. It does with me.

There is what I (and a great many others) regard as a plausible approach to the analysis of knowledge that leads quite naturally (not inevitably,³ but naturally) to a failure of closure. If knowledge that P requires one (or one's evidence) to exclude not all, but only all relevant, alternatives to P, then, it seems, one is committed to a failure of closure. The evidence that (by excluding all relevant alternatives) enables me to know there are cookies in the jar does not enable me to know that they (what I see in the jar) are not papier mâché fakes since papier mâché fakes are not (usually) a relevant alternative. So the evidence that gives me knowledge of P (there are cookies in the jar) can exist without evidence for knowing Q (they are not fake) despite my knowing that P implies Q. So closure fails.

Although this is, as I say, a natural line of reasoning, it isn't very effective against the skeptic if the only reason for embracing a relevant alternative analysis of knowledge is that it captures common-sense (basically antiskeptical) intuitions about when and what we know. If, that is, one's reasons for accepting a relevant alternatives analysis of knowledge is that it accords with our ordinary practice of claiming to know something (that there are cookies in the jar) without having specific evidence against possible mistakes (that, for example, they are not papier mâché fakes), then the argument against closure becomes too tightly circular to be effective against a skeptic. One uses premises – basically what is ordinarily regarded as the kind of evidence good enough to know – that no self-respecting skeptic would concede.

A dialectically more effective strategy is to provide an independent (of skepticism) analysis of knowledge that yields the result that only certain alternatives (to what is known) are evidentially relevant. This is the strategy adopted in Dretske (1969, 1970, 1971) and Nozick (1981). If knowledge is belief based on the kind of conclusive reasons I describe in Dretske (1971), for instance, then closure fails. Things turn out this way because one can have conclusive reasons, R, for P (R is a conclusive reason for P = $\text{d}_P R$ would not be true unless P were true) without having conclusive reasons for known consequences of P. For example,

- (C) It would not look to me as though there were cookies in the jar if there weren't cookies in the jar.

can be true – thus making the look of the jar (roughly: my experience of the jar) a conclusive reason for believing there are cookies in it – while

(C*) It would not look to me as though there were cookies in the jar if there were cleverly made papier mâché cookies in the jar.

is false – thus preventing the experience from being a conclusive reason for believing what I know to be implied by there being cookies in the jar – that they are not merely papier mâché fakes in the jar.

Unfortunately, though, the world has not rushed to embrace this theory of knowledge (or its information-theoretic successor in Dretske, 1981). Nonetheless, I think many of the arguments for something like a counterfactual condition on knowledge are good arguments, and I have yet to see effective counterarguments.⁴ I won't, however, recapitulate the history here. One doesn't make reasons better by repeating them. Suffice it to say that if one is sympathetic to something like (C) as a condition on seeing (hence, knowing) that there are cookies in the jar, one has taken a critical step towards denying closure. One has accepted a condition on knowledge that need not be, and often isn't, satisfied by known consequences of what is known.

3 Heavyweight Implications

I argued in section 1 that some ways – the perceptual ways – of coming to know P do not transmit evidential grounding for P to all the known implications of P. Ordinary things we come to know by perception always have heavyweight implications that are out of range: we cannot see (hear, smell, or feel) that they are true. If closure is not to result in complete skepticism about the external world, then, it seems that there must always be other ways of knowing that these heavyweight implications are true. I can see that there are cookies in the jar, but I cannot see that there is an external world. So if we have to know there is an external world in order to see there are cookies in a jar, as closure tells us we must, then there must be other (than visual) ways we know that there is an external world. There must be a way other than perception in which I know that I'm not a brain in a vat, that I'm not being massively deceived, that solipsism is false, that it is not all just a dream. What might these other ways of knowing be?

It is hard to see what other ways there could be since every way of knowing either fails to reach these heavyweight implications or generates its own heavyweight implications. *No* evidence⁵ transmits to *all* the implications of what it is evidence for.

Testimony, for instance, is considered to be an important source of knowledge that is distinct from perception. I didn't see the tire for myself, but I know it is flat because she told me it was flat. That is how I know. Yet, although I know that tires are material objects, I can learn, by being told, that I have a flat tire without thereby learning that solipsism is false – that there is something else in the universe besides me and my own ideas. No one, as I recall, ever told me there was something else besides me in the universe, but, if someone (a philosophy professor?) did, it is hard to believe that *that* is the way I learned there are other minds and a material world. Yet being

told is the way I come to know things – a great many things – that I know imply there is a material world. That there is a material world is something we know to be implied by what we are told (flat tires *are* material objects independent of me and my own thoughts and experiences) but it is not something we learn by being told. So, if closure holds, testimony – and, remember, this includes not only books, newspapers, and television but virtually our entire educational system – cannot be a source of knowledge unless we have some *other* way of finding out there is a material world. Some *other* way of discovering that solipsism is false.

Or consider memory. Although memory is not generally considered a way of *coming* to know (it is more like a way of *preserving* knowledge acquired by other means) it, too, has its heavyweight implications. Whatever cognitive mechanisms enable a person to preserve, by memory, the knowledge that P do not enable the person to preserve a knowledge of P's known consequences. If you remember having granola for breakfast this morning, then you know you had granola for breakfast this morning.⁶ If you remember going to the bank before stopping at the bakery, then you know you went to the bank before stopping at the bakery. But, as we all know, we couldn't have had granola for breakfast this morning, or stopped at the bank before going to the bakery, unless time, the succession of events, and, in particular, the past is real. That the past is real, however, is not something one remembers. It doesn't sound like something one *can* remember. This implication is more like a presupposition – something one simply takes for granted – in remembering the nitty gritty details of personal history – that one had granola for breakfast or went to the bank before going to the bakery. For most of the things we remember (*most* but perhaps not all: we sometimes remember what we are supposed *to do*) the reality of the past is a heavyweight implication. It is something we know to be implied by things we remember, but it is not itself something we remember. Our way of preserving knowledge is not a way of preserving what we know to be implied by what we preserve. If closure is true, though, we cannot remember what we had for breakfast this morning unless we know the past is real. Since memory isn't a way of finding out the past is real, what, then, is the method we have for finding this out? Perception? Testimony? Fossils? I don't think so. Neither did Bertrand Russell when he wondered how we could know the world wasn't created minutes ago complete with history books, fossils, and memory traces. If there is a past, fossils, memory, and history books tell us what happened in it – the historical details, as it were – but they cannot tell us what is implied by the existence of these details: that there actually was a past.

Instead of trying to catalogue all the different ways of coming to know – probably a hopeless task, anyway – let me try approaching this topic from a more general standpoint. Generally speaking, when we come to know that P is true, there is always some indicator, some source of information, that “tells us” P is true. In the case of direct perception, this is an experience of the very condition “P” describes. We come to know there are cookies in the jar by seeing (experiencing) the cookies themselves. In testimony the information arrives by more indirect means: friends, teachers, radio and television newscasters, witnesses, books, newspapers, and magazines act as intermediaries. The information arrives in verbal form: someone tells me – perhaps by scribbling a note – that so-and-so happened. We also rely on natural signs – tracks in the snow, rings in the stump, cloud formations, facial expressions, and

spontaneous behavior (yawns, sneezes, etc.). Photographs and measuring instruments indicate, they carry information about, conditions we do not – sometimes cannot – directly perceive. In each case, our way of coming to know that P always seems to depend on there being something that, by its revealing properties, indicates that P is true. There is always something that brings this information to us. It is important, therefore, to understand that the relationships we describe by speaking of one thing *indicating* or *carrying information* about another condition, P, are relations that do not relate this thing to (all) the known implications of P. Though we rely on a measuring instrument to indicate the value of Q – that it is, say, 5 units – the measuring instrument does not indicate, does not carry information, that Q is not being misrepresented. The position of a pointer (pointing at “5”) on a well functioning meter indicates that the value of Q is 5 units, but it does not thereby indicate that the instrument is not broken, not malfunctioning, not misrepresenting a value of 2 units (say) as 5 units. Even when instruments (and this includes the human senses) are in perfect working order, they do not – they cannot – carry information that what they are providing is genuine information and not misinformation. That isn’t an instrument’s job. The job of a Q-meter is to tell us about Q, not about its reliability in providing information about Q. So even though we know that if Q really is 5 (as the meter says), then Q can’t be 2 being misrepresented (by the meter) as 5, the meter gives us the first piece of information (tells us that Q is 5) without giving us the second piece of information: that Q isn’t really 2 being misrepresented as 5. That isn’t information the Q-meter supplies. It is a heavyweight implication of the information a Q-meter supplies. Anyone who really wanted to know whether the instrument was misrepresenting Q (when it registered “5”) should not look more carefully at the pointer. A pointer pointing at “5” does not carry *that* information. The information it carries (when things are working right) is that Q is 5, not that Q doesn’t just appear to be 5.

This is true of all indicators, all sources of information. That is why there is nothing in the world – either mental or material – that indicates that there is a material world. Nothing in the present that indicates there is a past. Skeptics, of course, have known this for a long time. Indicators carry information about cookies in jars, wine in bottles, and zebras in pens, not information about the heavyweight implications of these facts. That is why fuel gauges indicate, and thereby tell you, that you have gas in your tank without telling you that there is a material world (something implied by there being gas and gas tanks). That is why rings in the stump can indicate the age of a tree without indicating that the past is real. That is why birth certificates can provide information about when someone was born without providing information that the world (complete with birth certificates) was not created this morning. If there is an answer to skepticism about the past or about the material world, the answer will not be found by examining birth certificates or looking at tree stumps. If there is a past, if that heavyweight implication (of a tree being 200 years old or a person being 84 years old) is true, tree rings and birth certificates will not tell us it is true. That is something we take for granted *in*, not something we find out *by*, counting tree rings and reading birth certificates.

Insofar, then, as any accredited way of knowing P depends on information delivered about P, I conclude that *none* of our accredited ways of knowing about our mate-

rial world are capable of telling us that there is a material world, none of the accredited ways of finding out about what people feel and think are ways of finding out that they are not mindless zombies, and none of the accepted ways of finding out what, specifically, happened yesterday are ways of finding out that there was a yesterday. If this is true, it leaves us with a choice. Either

- (1) *Skepticism*: we retain closure and accept skepticism. I cannot know P, for any P, unless I know all the (known) implications of P. But I can't know all these. There is nothing I am aware of that indicates (carries information) that the heavyweight implications are true. So (given closure) we don't know any of the things – and this includes most things – that imply these heavyweight truths.

or

- (2) *Denial of closure*: we preserve ordinary knowledge by denying closure. We don't have to know the heavyweight implications of P to know P. I don't have to know I am not a deluded brain in a vat to see (hence, know) that there are cookies in the jar. I can, and often do, get information about P sufficient unto knowing that P is true without getting information about the heavyweight implications of P.

For anyone who thinks they know something about the material world, the choice is easy. Closure must be denied. There are some (known) implications of what we know – the heavyweight implications – that we do not have to know to be true in order to know to be true what implies them.

There is, though, another possibility. The choice between (1) and (2) is so painful that some philosophers prefer to try wriggling between the horns of this dilemma. Perhaps we can keep closure *and* have knowledge if we are willing to extend – by an act of courtesy, as it were – knowledge to heavyweight implications. Though I typically have no evidence that I am not being systematically deceived by the facts on which I base my belief that P,⁷ it might still be argued that our knowledge of heavyweight implications is special. Even though none of our accredited ways of knowing transmits evidence to them, we nonetheless get to know them anyway. These are things we get to know without having reasons to believe them, without being able to eliminate the possibilities (e.g. that I *am* being deceived) they directly compete with. Though I can't know they are zebras (not elephants or giraffes) in the pen without being able to distinguish (at this distance, in these circumstances) zebras from other sorts of animals, I can (according to the present hypothesis) know that they are not painted mules even though I can't (at this distance, in these circumstances) distinguish zebras from painted mules.

This strikes me as, at best, a bit of verbal hocus pocus designed to avoid skepticism while retaining closure. We have our cake and eat it too by simply stipulating that we know whatever we know to be implied by what we know even though we have no identifiable way of knowing it. In rejecting this option I assume, of course, that the visual appearance of a zebra – the sort of perceptual condition that normally leads one to say one can see that it is a zebra – is not evidence, is not itself a reason to think, that the zoo authorities did not replace the zebras with mules painted to

look like zebras. I assume that the appearance of cookies in a jar, the sort of condition that prompts one to say one can see there are cookies in the jar, is not itself a reason to believe that the experience of the cookies is not misleading or delusory in some way.⁸ Given these assumptions, it seems to me preposterous – abominable, if you will – to insist that we nonetheless know we are not being deceived in these ways because we know that *not* being deceived in these ways is implied by what our experience normally leads us to believe – that there are zebras in the pen and cookies in the jar. Despite its philosophical credentials (Moore, 1959), that sounds like chutzpah, not philosophy, to me.

Notes

- 1 I hereafter drop the distracting reminder that S, if he believes Q, must base his belief in Q on the two facts he knows to be true: (1) P, and (2) If P then Q.
- 2 The parenthetical “to be so” is important. Skepticism is irrelevant to what *objects* I see, whether, for instance, I see cookies in the jar, wine in the bottle, and zebras at the zoo. What I don’t see if skepticism is true are *the facts* – the fact that there are cookies in the jar, wine in the bottle, and a zebra in the pen. One doesn’t have to see (hence, know) these facts to see the objects that figure in these facts. Perception of objects (events, states of affairs, conditions) is immune from skeptical challenge.
- 3 Stine (1976), Cohen (1988, 1999) and Lewis (1996), while adopting a relevant alternative account of knowledge, retain closure by embracing a fairly radical form of contextualism. DeRose (1995) gives an externalist form of contextualism and argues that it retains closure while avoiding skepticism. For a general survey and discussion of options see Brueckner (1985).
- 4 Except for Martin (1975), which forced me to adopt an information-theoretic formulation of “conclusive reasons.” See, also, Martin (1983), where the same example is used against Nozick (1981). An oft mentioned criticism is that the theory is subject to the sorts of counterexamples (unpublished) Saul Kripke gave of Robert Nozick’s (1981) theory. I do not think these examples work against my own theory even though there is a superficial similarity between Nozick’s and my accounts (we both use counterfactuals to express the required relations between knower and known). Since I am not relying on this theory to defend my rejection of closure, a detailed defense of the theory is out of place here. For the cognoscenti, though, I say this much. When there are, in the relevant neighborhood, fake barns but no fake *red* barns (so that something might look like a barn without being a barn, but nothing would look like a red barn without being a red barn), it turns out that on Nozick’s theory you can track (Nozick’s term) red barns (you wouldn’t believe it was a red barn unless it was a red barn) without tracking barns. Thus, you can know of a red barn you see that it is a red barn but not that it is a barn. This result is an embarrassment even for someone (like Nozick) who denies closure. The example, however, is not effective against a “conclusive reason” (or information-theoretic) style analysis since these theories are formulated in terms not of a *belief* tracking a condition, but of one’s reasons or evidence (the condition causing you to believe) tracking the condition. S knows it is a barn if that feature of the evidence causing S to believe it is a barn would not exist if it were not a barn. In the case of perception, if its looking like a red barn is what is causing S to believe it is a barn, then S has conclusive reasons to believe it is a barn: it would not look that way (like a red barn) unless it was a barn, and its looking like a red barn is what is causing S to believe it is a

barn. Hence, he knows it is a barn. If, on the other hand, it is merely the building looking like a barn that is causing S to believe it is a barn (its color being irrelevant to the causing of the belief) then S does not know it is a barn. His experience carries the information that it is a barn (since it looks like a red barn to him and its looking like a red barn carries the information that it is a barn), but that isn't the aspect of experience that is causing him to believe. Knowledge is *information-caused belief* and in this second case the information (that it is a barn) isn't causing the belief. All this should be evident to anyone who has thought about an example Alvin Goldman (1975) introduced years ago. Even if I mistakenly take wolves to be dogs so that my belief that x is a dog does not "track" dogs in my environment I can nonetheless know of a dachshund, seen at close range in broad daylight, that it is a dog. What is crucial to knowing the dachshund is a dog is that it has a distinctive look (it is, in this respect, like a red barn), a look that only (dachshund) dogs have. If it is this distinctive look that causes me to believe it is a dog, I know it is a dog no matter how confused I am about wolves – no matter how much my beliefs about dogs fail to "track" dogs (no matter how many false dog beliefs I have).

- 5 Except, possibly, logically conclusive evidence. I ignore this possibility since no one, I assume, thinks we have logically conclusive reasons for beliefs about the external world, the past, and other minds – the sorts of beliefs to which skepticism is typically directed.
- 6 One can, of course, *think* one remembers P without knowing P. One can even do it when P is false. But this is best described as *ostensible*, not genuine, memory. If one truly remembers P, P is not only true, one knows it is true.
- 7 I think it is begging the question to object by saying that deception is most improbable. Yes, it is most improbable *given what we take ourselves to know about ourselves and the world we live in*, but, given the hypothesis in question, what allows us to help ourselves to this knowledge (about ourselves and the world we live in) to estimate the probabilities? Even if we helped ourselves to this (disputed) knowledge, we would still be left with the question: is the estimated improbability enough to say we know the hypothesis is false? I know the probabilities of my winning the lottery are vanishingly small, but that doesn't mean I know I'm going to lose.
- 8 Some philosophers, I know, would not grant me these assumptions. Klein (1981, 1995), for instance, argues that despite having nothing to show in the way of specific evidence for thinking the zoo authorities have not put a disguised mule in the zebra pen, one can nonetheless know they didn't do it because one can see that it is a zebra and one knows that if it is a zebra the zoo authorities didn't do that. Knowing that not-Q is implied by what one knows is good enough to know not-Q even if one cannot distinguish not-Q (it isn't a disguised mule) from Q (it is a disguised mule) and, therefore, has nothing to show for one's belief that not-Q. In the end, it may come down to a matter of taste (I hope not), but this view is as hard for me to swallow as others tell me the rejection of closure is for them.

References

- Brueckner, A. (1985) Skepticism and epistemic closure. *Philosophical Topics*, 13, 89–117.
- BonJour, L. (1987) Nozick, externalism, and skepticism. In L. Luper-Foy (ed.), *The Possibility of Knowledge*. Lanham, MD: Rowman & Littlefield.
- Cohen, S. (1988) How to be a fallibilist. *Philosophical Perspectives*, 2, 91–123.
- Cohen, S. (1999) Contextualism, skepticism, and the structure of reasons. In J. Tomberlin (ed.), *Philosophical Perspectives*. Atascadero, CA: Ridgeview.

- Davies, M. (1998) Externalism, architecturalism, and epistemic warrant. In C. Wright, B. C. Smith and C. Macdonald (eds), *Knowing Our Own Minds*. Oxford: Oxford University Press.
- Davies, M. (2000) Externalism and armchair knowledge. In P. Boghossian and C. Peacocke (eds), *New Essays on the A Priori*. Oxford: Oxford University Press.
- DeRose, K. (1995) Solving the skeptical problem. *Philosophical Review*, 104(1), 1–52.
- Dretske, F. (1969) *Seeing and Knowing*. Chicago: University of Chicago Press.
- Dretske, F. (1970) Epistemic operators. *Journal of Philosophy*, 67, 1007–23.
- Dretske, F. (1971) Conclusive reasons. *Australasian Journal of Philosophy*, 49, 1–22.
- Dretske, F. (1981) *Knowledge and the Flow of Information*. Cambridge, MA: MIT Press/A Bradford Book.
- Feldman, R. (1999) Contextualism and skepticism. In J. Tomberlin (ed.), *Philosophical Perspectives*. Atascadero, CA: Ridgeview.
- Fumerton, R. (1987) Nozick's epistemology. In S. Luper-Foy (ed.), *The Possibility of Knowledge*. Lanham, MD: Rowman and Littlefield.
- Goldman, A. (1975) Discrimination and perceptual knowledge. *Journal of Philosophy*, 73, 771–91.
- Grice, P. (1967) William James lectures at Harvard.
- Klein, P. (1981) *Certainty: A Refutation of Skepticism*. Minneapolis: University of Minnesota Press.
- Klein, P. (1995) Skepticism and closure: why the evil demon argument fails. *Philosophical Topics*, 23(1), 213–36.
- Lewis, D. (1996) Elusive knowledge. *Australian Journal of Philosophy*, 74(4), 549–67.
- Martin, R. (1975) Empirically conclusive reasons and scepticism. *Philosophical Studies*, 28, 215–17.
- Martin, R. (1983) Tracking Nozick's sceptic: a better method. *Analysis*, 43, 28–33.
- Moore, G. E. (1959) A proof of the external world. In *Philosophical Papers*. London: George Allen & Unwin.
- Nozick, R. (1981) *Philosophical Explanations*. Cambridge, MA: Harvard University Press.
- Stine, G. C. (1976) Skepticism, relevant alternatives, and deductive closure. *Philosophical Studies*, 29, 249–61.
- Tomberlin, J. (ed.) (1999) *Philosophical Perspectives*. Atascadero, CA: Ridgeview.
- Wright, C. (2003) Some reflections on the acquisition of warrant by inference. In S. Nuccetelli (ed.), *New Essays on Semantic Externalism, Skepticism, and Self Knowledge*. Cambridge, MA: MIT Press/A Bradford Book.

The Case for Closure

John Hawthorne

Introduction

Fred Dretske's discussion provides us with a case against epistemic closure that is, on its face, rather compelling. In brief, he reasons as follows. (1) There are various "heavy-weight" propositions, such as that I am not a brain in a vat, that I cannot come to know through use of my perceptual capacities. (2) Nor can I come to know them a

priori. So (3) I can't know them at all. But (4) rampant skepticism is unacceptable. I can come to know all sorts of ordinary propositions, such as that I have a hand, that there are cookies in my cupboard, that I used to frequent the Harborne fish and chip shop, and so on. But (5) those ordinary propositions are known by me to entail various heavyweight propositions of the sort that (by (1) and (2)) can't be known by me. So (6) it is false that knowledge transmits over known entailments.

Those who deny skepticism but accept closure – of whom there are many – will have to explain how we know various “heavyweight” skeptical hypotheses to be false. As Dretske sees it, this leads us to indulge in “verbal hocus pocus,” twisting and contorting the concept of knowledge to fit the twin desiderata of closure and anti-skepticism. Others reject (4), opting for skepticism. But Dretske, reasonably enough, expects most of us – even closure-lovers – to agree that acceptance of wholesale skepticism is even more “abominable” than a denial of closure. It would appear, then, that we have three choices: (a) abominable skepticism; (b) verbal hocus pocus; (c) the abandonment of the initially plausible sounding closure principle. Dretske hopes that we will come away feeling that the sober option is to abandon closure. Despite its self-confessed costs, his position is a seductive one.

In the remarks that follow, I shall not attempt definitive resolution of the very troubling issues that Dretske's admirable discussion raises. If there were some easily accessible locus of reflective equilibrium in the vicinity, we would surely have reached it by now. But each of the resolutions to the trichotomy carries undeniable costs.

My task here is not, then, to defend a particular resolution to the kind of puzzle that motivates Dretske to abandon closure. Rather it is to question the merits of his own preferred, closure-denying package. Once the force of a suitable epistemic closure principle has been properly appreciated and the untoward consequences of Dretske's own views have been properly exposed, the costs of Dretske's preferred package can be seen to be quite considerable.

In the discussion below, I shall proceed as follows. Having in section 1 said a little about the closure principle I am most interested in, I shall in section 2 underscore the intuitive cost of abandoning that principle. In section 3, I shall point to some difficulties attending to the project of devising a restricted version of closure. In section 4, I shall make some brief remarks about the kind of puzzle that motivates Dretske to abandon closure.

1 Versions of Epistemic Closure

Epistemic closure is sometimes formulated as the principle (or, more accurately, schema) that

$$(1) (S \text{ knows } P \text{ and Necessarily } (P \supset Q)) \supset S \text{ knows } Q$$

But this simple formulation does not present a defensible principle: if the relevant conditional is unknown to the subject, we should hardly be surprised if he were to know its consequent though not its antecedent.¹

Recognizing this, others formulate epistemic closure as the principle that

Is Knowledge Closed under Known Entailment?

27

(2) (S knows P and S knows that Necessarily $(P \supset Q)$) \supset S knows Q.²

But that doesn't seem right either. If one knows P and the relevant conditional, one still might not know Q on account of not having performed the relevant deductive inference. Knowing the logical consequences of what one knows takes cognitive labor.

Dretske is well aware of all this, of course, and formulates a more guarded version of the closure principle:

(3) If one knows P and one knows that Necessarily $(P \supset Q)$, one has all that it takes, evidentially speaking, to know that Q.³

Here is the picture. Suppose one has a body of evidence such that, were one to base a belief on that evidence, one would thereby come to know P. Even if one doesn't so form a belief, it is nonetheless natural to say that one has all it takes, evidentially speaking, to know P.

But version (3) of the closure principle invites at least two sorts of objections other than those that concern Dretske in his discussion.⁴

The first flows from the thought that one cannot always know the logical consequences of the propositions that one knows on account of the fact that small risks add up to big risks. Such a thought flows naturally from the view that knowledge that P requires that the epistemic probability of P be above a certain threshold N (where N is less than 1). I myself am not especially attracted to this "threshold conception" of knowledge. But if one were to embrace such a conception, then (3) would only be acceptable under the somewhat unrealistic idealization that one knows all necessary truths with certainty. And this is unrealistic. Suppose someone who is quite reliable tells me that the atomic number of calcium is 20. Suppose I know for certain that that if calcium has atomic number 20 then it necessarily has that number. If my rational credence that he is telling the truth is less than one, then the same is intuitively so for the proposition

Necessarily (Calcium exists \supset Something with atomic number 20 exists)

So if one embraces the threshold conception of knowledge, (3) is unacceptable. This problem can be sidestepped by moving to (3')

(3') If one knows P and knows with *a priori* certainty that Necessarily $P \supset Q$, one has all that it takes, evidentially speaking, to know Q.

A second sort of worry, which confronts both (3) and (3') turns on some delicate issues about the relationship between logical consequence and rationality. Consider the following case, which builds on Saul Kripke's famous Padarevski example (see Kripke, 1979). Suppose I overhear someone (who is utterly trustworthy) say "Fred Padarevski is Bob Padarevski's brother" and I accept that sentence. I overhear someone else (who is utterly trustworthy) say "Ivan Padarevski is Fred Padarevski's father" and also accept that sentence. However, I do not realize that the token "Fred Padarevski" as used by the second individual corefers with the token of "Fred Padarevski" used by the first.

Having accepted the second sentence, I apply the necessity of origins principle and go on to accept “Necessarily if Fred Paderevski is Bob Paderevski’s brother, Ivan Paderevski is Bob Paderevski’s father.” Now many of us hold a conception of belief attribution according to which an onlooker could truly say “John knows that Bob Paderevski is Fred Paderevski’s brother and John knows that necessarily, if Bob Paderevski is Fred Paderevski’s brother, then Ivan Paderevski is Bob Paderevski’s father.”⁵ But it is quite clear that I am in no position to know that Ivan Paderevski is Bob Paderevski’s father. (I am agnostic, after all, as to whether the name “Fred Paderevski” as used by the first interlocutor is coreferential with the name “Fred Paderevski” as used by the second interlocutor.) And without getting fully clear on what “evidentially speaking” is supposed to amount to, it seems at least dubious whether I “have what it takes, evidentially speaking, to know that Ivan Paderevski is Bob Paderevski’s father.” There is plenty more to say here, but because the issues that primarily concern Dretske do not turn on these matters, I will leave things somewhat hanging. Suffice it to say there is a good case to be made that, even if one is ideally rational, one will often be in no position to recognize the logical consequences of one’s knowledge on account of diverse modes of linguistic representation of the propositions known.

If the closure debate were over (3) or (3’), then I would be the wrong recruit for defending closure. Fortunately, the issues just raised about Dretske’s preferred formulation are somewhat tangential to the core of the debate. Here is a version that avoids these issues:

- (4) 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.⁶

(Why “while retaining one’s knowledge that P”? Well, one might competently deduce Q but have gotten counterevidence to P by the time one comes to believe Q.)

The “Paderevski case” does not pose problems for (4), since in that case a competent deduction to the consequent of the conditional is, intuitively, not available. And (4) offers in any case a more satisfying development of the closure intuition. The core idea behind closure is that we can add to what we know by performing deductions on what we already know. The idea that knowing both P and Necessarily $P \supset Q$ is evidentially sufficient for knowing Q is, it seems to me, parasitic on the idea that *a priori* deduction from P yields knowledge of the propositions so deduced. My concern, then, will be to defend (4), rather than Dretske’s own stated target, (3). To do so is not to frame the issue unfairly. The concerns that Dretske raises about closure carry over perfectly well to (4), and most of what I shall say in response carries over perfectly well to Dretske’s criticisms of (3).

It is worth noting in passing that there is a more ambitious version of closure than (4), one that generalizes to deductions with more than one premise, namely:

- (5) If one knows some premises and competently deduces Q from those premises, thereby coming to believe Q, while retaining one’s knowledge of those premises throughout, one comes to know that Q.

Like (4), (5) has great prima facie intuitive plausibility. Indeed, (4) and (5) seem to originate from the core idea that one can add to what one knows by deductive inference from what one already knows. But defending (5) requires additional work. In particular, one has to confront the concern already mentioned, that “small risks add up to big risks,” a concern that has no force against (4).⁷ One has to also confront intuitive counterevidence arising from consideration of particular cases: allowing that a businessmen knows that he will meet X on Monday, Y on Tuesday and so forth, it still seems outrageous to allow that he knows some long conjunction of those claims, since the possibility that one of the people on the list pulls out sick seems not to be one that can be disregarded. So defending (5) requires, on the one hand, a conception of epistemic risk according to which knowledge involves no risk⁸ and, on the other, some satisfactory resolution of the counterintuitions generated by a consideration of cases. I will not undertake any such defense here.

2 In Favor of Closure Principle (4)

Dretske and I would *both* agree that (4) is intuitive. It is unquestionably odd to concede that someone knows some P, to recognize that the person went on to realize that some Q follows from P and to believe Q on that basis, but to deny that the person knows Q.⁹

Now as Dretske points out, there are claims that are “conversational abominations” but that are perfectly true nevertheless. The contrast between truth and conversational propriety is, indeed, a very important one. For example, it is odd to actually assert “I know I have a hand” in nearly all ordinary conversational contexts. But we should not let such facts about conversational impropriety induce us into questioning the truth of “I know I have a hand.” Might a parallel diagnosis be applied here? It seems not. In the case of “I know I have a hand,” one can produce a convincing account in terms of the conversational mechanisms that Grice described concerning why the relevant speech is true though infelicitous. Roughly, we can presume participants in the conversation already know that I know that I have a hand and that, in keeping with mutually accepted conversational maxims, they will reckon that I am trying to be informative in my utterances. But “I know I have a hand” would only be informative if there were some real question as to its truth. So if I utter “I know I have a hand” then I can only be reckoned a cooperative conversant by my interlocutors on the assumption that there was a real question as to whether I have a hand. If there is no real question, my utterance misleadingly suggests there is such a question – hence the infelicity in most contexts. But what sort of mechanism could explain why “He knew P and deduced Q from P but didn’t know Q” is true but conversationally infelicitous? No analogous Gricean narrative seems to be available here.

Moreover, the relevant infelicity intuitions arise not merely at the level of conversation but at the level of thought. If I think someone knows P and deduced Q from P, then I will not merely reckon it inappropriate to say “He still doesn’t know Q”: I will actually form the belief that he does know Q. Correlatively, when I reflect on principle (4), I form the belief that it is obviously correct. If an utterance u means P, then conversational mechanisms may explain away why we don’t utter u even though

it is correct, but they will not explain why we believe that which is expressed by the negation of *u*.

Consider by analogy the transitivity of the *taller than* relation. Suppose, somehow, that the taller than relation isn't transitive. If God told me such a thing, I would not conclude that I had been guilty merely of conversational abominations in making utterances that presupposed its transitivity. On the contrary, I would conclude that I had some seriously mistaken beliefs about the nature of the *taller than* relation. The closure principle seems to me like that. If Dretske is right, then I am prone to forming false beliefs – some involving conditionals, some not – about what others know by applying a closure principle that does not have full generality.

So our commitment to something like (4) cannot be explained away by appeal to pragmatic mechanisms. But let me be honest here. While principle (4) is compelling, it is not as *manifestly* obvious as the transitivity of the *taller than* relation. We would think it ridiculous if someone were to write a paper challenging the latter. And while we might say, by way of rhetoric, that Dretske is being ridiculous in denying (4) (or its counterpart (3)), the cases seem different. Perhaps, then, we should work to provide some supplementary considerations that reinforce our commitment to (4)'s intuitive status. Here are two lines of thought to that end.

Argument (A): closure and distribution

Begin with a principle that seems extremely compelling.

The Equivalence Principle: If one knows *a priori* (with certainty) that P is equivalent to Q and knows P, and competently deduces Q from P (retaining one's knowledge that P), one knows Q.

Interestingly, Dretske's reasons for denying closure have no force against the Equivalence Principle. His argument against closure relies on the following idea. Following recent usage, let us say that R is "sensitive" to P just in case were P not the case, R would not be the case. Suppose one believes P on the basis of R, and that P entails Q. R may be sensitive to P and still not to Q. But notice that where P and Q are equivalent, there can be no such basis for claiming that while R can underwrite knowledge that P, it cannot underwrite knowledge that Q. We may thus safely assume that Dretske will accept the Equivalence principle. Here is a second, equally compelling, principle:

Distribution: If one knows the conjunction of P and Q, then as long as one is able to deduce P, one is in a position to know that P (and as long as one is able to deduce Q, one is in a position to know that Q).

(Often a stronger principle is defended: that if one knows P and Q, one knows P and one knows Q. I do not need that here, though I do not wish to question it here either.)

Distribution seems incredibly plausible. How could one know that P and Q but not be in a position to know that P by deduction? Dretske, however, is committed to denying distribution. Suppose one knows some glass *g* is full of wine on the basis of

perception (coupled, perhaps, with various background beliefs). The proposition that g is full of wine is *a priori* equivalent to the proposition

g is full of wine and $\sim g$ is full of non-wine that is colored like wine.

So by equivalence one knows that conjunction. Supposing distribution, one is in a position to know that

$\sim g$ is full of non-wine that is colored like wine.

But the whole point of Dretske's position is to deny that one can know the latter in the type of situation that we have in mind.

What does this show? I do not offer the argument as a decisive refutation of Dretske's position. What we learn is that in giving up closure we are forced to give up another extremely intuitive and compelling principle, namely the distribution principle. If we are to play the game of counting costs, it must surely be reckoned a significant cost of Dretske's view that he must eschew Distribution.¹⁰

Argument (B): knowledge and assertion

An idea with growing popularity among epistemologists is that there is a fundamental tie between knowledge and the practice of asserting (see Unger, 1975; De Rose, 1996; Williamson, 2000). Of course, we sometimes assert that which we do not know. But it seems that when we do, we at least represent ourselves as knowing, a fact that is in turn explicable by Timothy Williamson's conjecture that it is the fundamental norm of assertion that one ought not to assert that which one does not know. Combine this thesis with Dretske's view and additional oddities arise. Suppose Q is a "heavyweight" consequence of P and S knows P and also that P entails Q . 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 (see Carroll, 1895). If we embrace the thesis that knowledge is the norm of assertion, then a Dretske style position on closure will turn us into a facsimile of Carroll's Tortoise.

3 Heavyweight Propositions and Conclusive Reasons

Dretske does not deny that deductive inference often succeeds in adding to our stock of knowledge. But he restricts its knowledge-extending scope: deductive inference adds to our knowledge in just those cases where the deductively reached conclusion is not a "heavyweight" proposition. *Prima facie*, the picture is quite attractive. Most

deductive extensions of our belief set involve no heavyweight propositions. And in those cases, Dretske allows, the conclusions can be known as long as the premises are too. Though closure fails in full generality, it holds as a *ceteris paribus* generalization. Admittedly, in those special cases where we do arrive at heavyweight propositions, deductive inference does not extend knowledge. But it is not as if we are ordinarily misled here. On the contrary, inner alarm bells go off. If, having formed the belief that she is sitting, an ordinary person is invited to infer that she is not a brain in a vat (with an accompanying explanation of what that hypothesis involves), she may well not embrace that conclusion, even though it is a deductive consequence of what she just claimed. On the contrary, she will likely say, "Gee, I can't know that." It begins to look as if our actual practice conforms fairly well to what Dretske's principles would predict.

But matters become more complicated when we explore the concept of a "heavyweight" proposition. Let us begin with a somewhat vague, but nevertheless workable, gloss on the crucial notion. Let P be a "heavyweight proposition" just in case we all have some strong inclination to think that P is not the sort of thing that one can know by the exercise of reason alone and also that P is not the sort of thing that one can know by use of one's perceptual faculties (even aided by reason). Rough and ready as the characterization is, it captures a familiar category. And Dretske is just right when he points out that we all have some strong inclination to think that we cannot know by perception that we are not looking at a white thing illuminated by a red light (when in fact something looks red), that we are not brains in vats, that we are not looking at a cleverly disguised mule (instead of a zebra), and so on. And he is right when he points out that we all have a strong inclination to say that we cannot know such propositions by rational reflection either. In conceding this, I am not claiming that we do not know such things. But I am conceding that I have an immediate and strong inclination to say that I do not know such things in a way that I do not, for example, have any such immediate inclination to say that I do not know that I am working on a computer right now. Of course, I might derivatively generate some inclination to say that I do not know the latter via reflection on the former. But in the case of the former, doubts are natural and immediate, whereas in the case of the latter, doubts tend to be derivative. We should all at least agree that the relevant contrast is a salient fact about human psychology and that it is one of the bases of epistemic puzzlement. Let us call those propositions that generate immediate epistemic doubt "manifestly heavyweight propositions."

In response to the phenomenon of manifestly heavyweight propositions, one familiar reaction is to combat the inclinations towards doubt. Another is to allow doubt to spread to more humdrum propositions so that skepticism prevails. A third less familiar response is Dretske's – to reject closure and to offer a theoretical characterization of those propositions for which closure fails.

But the problem with this sort of approach is that the set of knowledge-eluding propositions identified by the theoretical characterization turns out to be markedly different from the set of manifestly heavyweight propositions that motivated the approach in the first place. Consider, by way of example, Robert Nozick's approach, in its simplest version. Nozick is initially taken with the idea that a necessary condition for knowing P is the following:

Is Knowledge Closed under Known Entailment?

33

If P were not the case, S wouldn't believe P.

But such a principle will rule out knowledge of many non-heavyweight propositions as well. Suppose a real dog and a fake cat are in a room, the former keeping the latter from view. I look at the dog and form the belief that there is a dog in the room. From this, I infer there is an animal in the room. Suppose further that if there hadn't been a dog in the room, I would have seen the fake cat and formed the belief that it was a (real) animal. Then my belief that there is a dog in the room passes the Nozick test but my belief that there is an animal in the room does not. But my belief that there is an animal in the room in no way involves a heavyweight proposition. Examples such as this one put considerable pressure on closure-challengers like Nozick to "fix" their accounts of the necessary conditions on knowledge.¹¹ The aim of the game is to provide an account that precludes inference from what is known to manifestly heavyweight propositions but that allows inference from what is known to those that are not. If I point out that by the theorist's lights I can know that I have hands but not that I am not a brain in a vat, the theorist will welcome the result. But if I point out that by the theorist's lights there are situations where I can know that there is a cat in front of me but not that there is an animal in front of me, the theorist will hardly celebrate this result. On the contrary, he will scurry back to the epistemological laboratory to contrive an account that delivers the welcome result while avoiding the embarrassing one. Time and again, closure-deniers fail to pull this off. I believe that Dretske's own preferred account of knowledge offers an instance of this pattern of failure.

To begin, let us get clear about the form of his account. Dretske begins with the idea that to know that P one has to "rule out" or "exclude" all relevant alternatives to P. He claims in passing that this approach to the analysis of knowledge leads "quite naturally (not inevitably, but naturally)" to a denial of closure. I find this claim hard to evaluate without a clear sense of what "ruling out" or "excluding" amounts to. One very natural gloss on "ruling out P" is "knowing not-P." On this gloss, the suggestion is that for each P that one knows at t, there are some propositions incompatible with P the falsehood of which one has to know at t as well. Whether or not this is true – though I doubt that it is – I doubt that this is the intended meaning. Perhaps "rule out" means "to be able to know not-P" or "have what it takes, evidentially speaking, to know not-P." Assuming suitable deductive prowess, the suggestion that knowing P entails one can rule out relevant alternatives to P hardly encodes a challenge to closure. Indeed, assuming closure, one would, on this construal, be encouraged to think that if one knows P one can know the falsehood of any old alternative to P, relevant or not.

Fortunately, Dretske offers us a positive constraint on knowledge that bypasses this issue and that does indeed yield the result that closure fails.¹² The constraint is based on the idea of "conclusive reasons." Suppose I believe P on the basis of reasons R. Dretske tells us that if I know P, then I would not have reasons R unless P.

Let us say that reasons R are conclusive for P just in case this condition is satisfied. Suppose it is a necessary condition for knowledge that one's reasons are conclusive. Then, assuming we are not skeptics, and assuming a very natural gloss on one's reasons and the relevant counterfactuals, closure does indeed fail. Suppose A

believes there is a zebra in the cage (and that this is the only animal in the cage) and A's reasons e are a set of zebra-ish experiences (coupled, perhaps, with some background beliefs). It is natural to suppose that if there weren't a zebra there, A wouldn't have had e (and if there were other animals there as well, A wouldn't have had e (either). So A's reasons are conclusive. Let us assume further that A knows on some occasion that the only animal in the cage is a zebra, her reasons being roughly as just described.¹³ (This is, of course, an additional assumption. The conclusive reasons principle, as we have formulated it, is only a necessary condition on knowledge.¹⁴) A then deduces and comes to believe that there is no non-zebra in the cage cleverly disguised as a zebra. Let us suppose, as is natural, A's reasons are one or both of e and her belief that the only animal in the cage is a zebra. But if there were a cleverly disguised mule in the cage, she might well still have believed there was a zebra in the cage and that it was the only animal. So A does not have conclusive reasons for her belief that there is no cleverly disguised mule in the cage.¹⁵ The necessary condition is not satisfied and so she does not know.¹⁶

Note further an obvious advantage of the Dretske approach over the simple version of the Nozick approach, one that Dretske is fully aware of.¹⁷ Return to our "fake cat" case (see Dretske, note 4). One would not wish to concede in that case that one knows that there is a dog in the room but does not know that there is an animal in the room. But Dretske is not forced to such a concession. Suppose one's reasons for believing that there is an animal in the room are one or both of: (a) doggish experiences; (b) the belief that there is a dog in the room. If there were no animal in the room one would not have had *those* reasons for believing there to be an animal in the room (though one might have had others). So one turns out to have conclusive reasons for believing that there is an animal in the room. No embarrassing concession is mandated.

Nevertheless, embarrassment is close at hand. It seems evident that Dretske's account is designed to deliver the conclusion that one knows ordinary propositions and their non-manifestly heavyweight consequences, while remaining ignorant of their manifestly heavyweight consequences. But it does not come close to delivering that result. It turns out that we all too often have conclusive reasons for manifestly heavyweight propositions (in which case Dretske has not in fact provided an effective barrier to knowing such propositions), and all too often lack conclusive reasons for *a priori* consequences of known propositions, even though those consequences are not manifestly heavyweight. In these cases, not surprisingly, it is embarrassing to stick to one's theoretic guns and deny that the consequences are known.

Category 1: manifestly heavyweight with conclusive reasons

Case 1. As the discussion above in effect revealed, there are many conjunctions that are manifestly heavyweight but for which I have conclusive reasons. Thus, while I might lack conclusive reasons for the proposition \sim I am a brain in a vat, I will (supposing I have a headache) have conclusive reasons for I have a headache and \sim I am a brain in a vat. My reasons for that conjunction include my headache. Were the conjunction false, I would not, then, have had my reasons. (The closest worlds where that conjunction is false are ones where I do not have a headache. Stated without recourse

to possible worlds language: were that conjunction false, it would be because I didn't have a headache.)

My putative knowledge of that conjunction cannot, then, be challenged on the grounds that I lack conclusive reasons. But the proposition that I am credited with knowing is every bit as apt to raise inner alarm bells as the proposition that ~I am a brain in a vat and will thus come out as manifestly heavyweight.

Case 2. I go to the zoo and see a bird flying around in a cage. I form the belief that I am not looking at an inanimate object cleverly disguised to look like an animate object. That is intuitively a manifestly heavyweight proposition. But suppose (as seems reasonable) that while it is just about possible to make an inanimate object appear like a flying bird, there are much easier ways to make an inanimate object look like an animate object. For example, it is much easier to pull off the task of creating an inanimate object that looks like a turtle than to pull off the task of creating an inanimate object that looks like a flying bird. Suppose then that "at the closest worlds" where there is an inanimate object in the cage, there is nothing that looks like a bird in the cage: instead there is something that looks like a far more sedate animal. Suppose, as seems plausible, my reason for thinking that there I am not looking at a cleverly disguised inanimate object in the cage is that I have experiences as of a bird flying around. Then I have conclusive reasons, in Dretske's sense, for the proposition that there is no cleverly disguised inanimate object.

Case 3. I see a cookie. I form the belief that there is a mind-independent object that is roughly five feet in front of me. My sense is that Dretske wishes to classify such high-filutin' theoretically loaded beliefs as "heavyweight." But that one does not come out as heavyweight by the conclusive reasons test. For presumably the closest worlds where there is no mind-independent physical object five feet in front of me are not ones where some bizarre metaphysics holds but, rather, worlds with laws like this one where there is no physical object at all in front of me. Alternatively put: if there weren't a mind-independent object about five feet in front of me, that would be because that region of space was unoccupied. But if that had happened, then (in the normal case) I wouldn't have had the experiences as of a cookie five feet away from me.

Category 2: not manifestly heavyweight but lacking conclusive reasons

Case 1. I eat some salmon for dinner. I am no glutton. I eat a modest quantity and form the belief that I have eaten less than one pound of salmon. I infer that I have eaten less than 14 pounds of salmon. In fact my perceptual system is very reliable indeed. In those nearby possible worlds where I feast on salmon and eat, say, a pound and a half, I do not believe that I have eaten less than a pound. Dretske, no skeptic, will happily concede that I know that I have eaten less than a pound of salmon. Suppose (and for all I know this is correct) that while it is utterly unlikely that any human being would eat over 14 pounds of salmon, doing this would induce (among other things) severe hallucinations. Indeed, it might even induce the hallucination that one had eaten a rather small quantity of salmon.¹⁸ Thus it might be true that if, God

forbid, I had eaten over 14 pounds of salmon, I may have had the eaten-less-than-one-pound-ish visual experiences that I have at the actual world. Thus while it may be true that I would not have had my reasons R for believing that I have eaten less than a pound of salmon unless I had, it is not true that I would not have had my reasons R for believing that I have eaten less than 14 pounds of salmon unless I had. The apparatus (in conjunction with the verdict that I know, in the given case, that I have eaten less than a pound of salmon) delivers the embarrassing result that I know that I have eaten less than a pound of salmon but not that I have eaten less than 14 pounds (even though I believe the latter and have deduced it from the former).

Case 2. I think on Monday that I am going to meet you on Wednesday on the basis of familiar sorts of reasons. (Fill them out in a natural way.) I meet you on Wednesday. It is natural to say, in many such cases, that my belief on Monday was a piece of knowledge. Consider now the counterfactual: if you hadn't met me on Wednesday, I would not have had those reasons on Monday for thinking that you would meet me on Wednesday. I do not find that counterfactual very compelling. (I have a strong inclination not to "backtrack" when evaluating the counterfactual – in effect accepting David Lewis's (1973) advice to hold the past relative to time *t* fixed when evaluating counterfactuals concerning *t*.) So my reasons turn out not to be conclusive, since they fail Dretske's counterfactual test.¹⁹

(There is an additional worry in the vicinity: counterfactuals are slippery and our modes of evaluating them very shifty. Dangerous, then, to build an account of knowledge on them unless one is happy for the truth of knowledge claims to slip and slide with the varying similarity metrics that may guide verdicts about counterfactuals. Let me illustrate with a case. Suppose, unknown to me, the beach is closed when the authorities believe the neighboring waters to be shark-infested. I go to the beach and ask "Is the water shark-infested?" The lifeguard says "no." Now if the lifeguard had believed the waters to be shark-infested, he wouldn't even have been on duty that day – the beach would have been closed. In that case I wouldn't have even gotten to ask him that question. Now consider the counterfactual "If the waters had been shark-infested, the lifeguard wouldn't have said that they weren't." Our evaluation fluctuates according to what we hold fixed. If we hold fixed that I am talking to a lifeguard, then, intuitively, the counterfactual comes out false. Meanwhile, if we are allowed to backtrack, imagining that I wouldn't have even reached the beach if the waters had been shark-infested, the counterfactual comes out true. Yet the parameter on which the truth of the counterfactual depends seems not to be one upon which the truth of the knowledge claim "I now know that the water is not shark-infested" similarly depends.)

The cumulative lesson is clear enough. Dretske's machinery is intended to align itself with our instinctive verdicts about what we can and cannot know by perception and reason. But, quite simply, it draws the can/cannot line in a very different place than intuition does. Of course, one might retreat to an account that gets more intuitively satisfying results by brute force, along the lines of:

If one knows *P* and deduces *Q* from *P* then one knows *Q*, unless *Q* is a manifestly heavy-weight proposition.

But we should all agree that this would be rather unsatisfying. But it remains an open question whether any reasonably satisfying non *ad hoc* restriction on closure is available. Epistemologists far too often assume that a discrepancy in our epistemic intuitions about pairs of cases can be accounted for in terms of interesting structural differences between the cases – differences in the pattern of counterfactuals and so on – without paying due attention to the myriad ways in which our perceptions of epistemic risk are at the mercy of incidental features of the way that the case is framed, of distortions induced by particularly vivid descriptions of this or that source of error, and so on (see Kahneman et al., 1982, 2000). I myself hold out very little hope for a plausible restriction on closure that allows with the skeptic that we are in no position to know we are not brains in vats but that allows deduction to extend knowledge in the normal case.

I close this section with one final intuition pump. Child A visits the zoo and is told by his trustworthy parent that he will get a prize if the first animal he sees is a zebra. The parent adds that the child may get a prize anyway. But whatever else, the child will definitely get a prize if the cage at the zoo entrance contains a zebra. Child B visits the zoo and is told the same thing by his parent. The parents of A and B consult. Parent A is both a philosopher and very generous. He has in fact decided that he will give child A a prize just as long as the cage at the zoo entrance does not contain a mule cleverly disguised as a zebra. (He resolves to check with the honest zookeeper before dispensing the prize.) The parent of B is less generous and less imaginative. He resolves to give child B a prize iff the first animal is either a zebra or a lion. The parents consult. Since they are walking together, they realize that if the second child wins a prize the first will also – that, indeed, it is easier for A to win a prize than for B to win a prize. They tell this to the children (though they do not explain in detail why that discrepancy obtains). The parents are very trustworthy, transmitting knowledge by testimony. Thus child A knows (a) if the cage at the zoo entrance contains a zebra, he will get a prize, (b) if the cage at the zoo entrance contains a zebra, B will get a prize, and (c) if child B gets a prize, child A will get a prize. Child B knows (a) that if the cage at the zoo entrance contains a zebra, he will get a prize, (b) that if the animal is a zebra, A will also get a prize, and also (c) that if he, B, gets a prize, child A will get a prize.

They see a zebra. Child A knows there is a zebra there. Child B knows there is a zebra there. Child A infers (a) that he will get a prize and (b) that B will get a prize. Child B infers (a) that he will get a prize and (b) that A will get a prize. Apply the Dretske machinery and we get some curious results. A does not know he will get a prize, it would seem. For if A weren't to get a prize it would be because he saw a cleverly disguised mule, in which case A would still think he was going to get a prize. But A knows B will get a prize! Furthermore, A knows that it is easier for A to win a prize than for B to – and in particular, that if B gets a prize, A will get a prize. B, meanwhile, does know that he will get a prize and knows that if he will get a prize, then so will A – since he also knows that the rules of the game are that A is guaranteed a prize if B wins. But, by parity of reasoning, B does not know that A will get a prize. At best, these conclusions are a bit dizzying.

4 Residual Puzzles

Suppose we do not embrace Dretske's rejection of closure. We are still left with the puzzle that led him to the rejection in the first place. Let us focus on an example. Suppose I learn from a reputable newspaper that Manchester United beat Coventry City. That piece of information, in combination with the fact that the newspaper says that Manchester United beat Coventry City, entails that

~Due to a misprint, the newspaper said that Manchester United beat Coventry City when in fact Manchester United lost to Coventry City.

Call this proposition "Misprint." On the one hand, it seems odd to suppose that one can know Misprint in advance of reading the newspaper; on the other, it seems odd to suppose that one can come to know Misprint by looking at the newspaper.²⁰

One prima facie attractive way out is that of the contextualist. Roughly, the contextualist says that the standards for knowledge ascription vary according to what the ascriber is worried about. If one worries about misprints, one's standards for knowledge ascriptions in testimony cases go up, and that is why we balk at *saying*, in the case described, that one knows Misprint upon looking at the newspaper. But the contextualist saves closure: in contexts where such worries are salient, the ascription "He knows the Manchester United result" is also false. So we are not left ascribing knowledge to the premises of a deductive argument while denying knowledge of the conclusion.

I am by no means convinced that contextualism holds the key to Dretske's puzzle. But Dretske's own way with contextualism is too quick, betraying a flawed understanding of that approach. He attributes to the contextualist the thesis that reasonably experienced adults know such things as that there are cookies only as long as they do not seriously consider the heavyweight implications of such beliefs. This is wrong. Contextualism says that the standards of "know" are *ascriber* sensitive. Suppose that A isn't considering heavyweight possibilities but that B is. In particular suppose that B believes P and also goes on to believe some heavyweight consequence of P by deduction. Suppose A says (1) "B knows P" and, moreover, (2) "B knows anything he has deductively inferred from P and thereby come to believe." If B is in a context where "A knows P," in her mouth, expresses a truth, then she will be in a context where (2), in her mouth, expresses a truth as well. That the *subject* of a knowledge ascription attends to heavyweight propositions does not undermine knowledge ascriptions by an ascriber who does not worry about this or that skeptical scenario. Indeed, in the case just described, the standard contextualist will allow that the extension of "knows" in the mouth of A may include the pair $\langle B, Q \rangle$, where Q is a heavyweight proposition. Relatedly, it is simply not correct that contextualism delivers the result that "philosophers are the most ignorant people in the world."

Contextualism remains one prima facie promising avenue for resolving Dretske's puzzle. Are there others? Without pretending to be able to resolve perennial problems in a few short paragraphs, let me at least gesture at one other move that is perfectly consistent with closure.

Begin with the humble phenomenon of epistemic anxiety, cases where human beings really begin to have genuine doubts about what they believe. Suppose, for example, I worry that my partner is going to leave me or that the plane I am on is going to crash. In such situations, I may still have the betting inclinations symptomatic of a believer. For example, despite my worries, I may be willing to bet heavily on the plane's safe arrival. But it is an obvious feature of the way the concept of knowledge works that in such a situation one is not inclined to self-ascribe knowledge of the proposition that is the target of one's anxiety. This might be accounted for variously. Perhaps knowledge requires a kind of conviction that is incompatible with such anxiety. Or perhaps instead, while knowledge is in fact compatible with such anxiety, the latter is constituted in part by a belief that one doesn't know.²¹

Notice that such anxiety, once generated and unalleviated – either by informal therapy or by a simple shift of attention – tends to spread. If one has it about Q and P entails Q, one acquires it with respect to P (when the question arises). Someone might know that his partner will never leave him and this may entail that she will never leave him for Mr X. But if he runs through the relevant inference, that may induce anxiety about both the inferred proposition and the original (we all know how to fill in the details!)

On such a picture, there are two sorts of situations. On the one hand, there are standard cases where deductive inference adds to our stock of knowledge. On the other, there are situations where deductive inference brings to light sources of anxiety that may, in some way or other, undercut the epistemic status – or at least the believed epistemic status – of one's premises. An account of the human epistemic condition should give due place to both mechanisms. But it remains unclear to me whether a denial of epistemic closure would have any useful role to play in any such an account.²² On the contrary, the reality of closure explains the epistemic goodness of deductive inference as well as why epistemic anxiety tends to spread.

But still, what of the puzzle itself? Do I know Misprint in advance? And then, by parity of reasoning, can I have contingent *a priori* knowledge that one is not a brain in a vat? Or is it that, as long as anxiety can be kept at bay, we can know the latter only by perception in combination with deductive inference (intuitions and Bayesian style arguments to the contrary notwithstanding²³)? Better that we appreciate (with Dretske) the force of this problem than rush too quickly to try to solve it.²⁴

Notes

- 1 There are those who adopt a conception of thought according to which the thought that P is identical to the thought that Q just in case P and Q are true in the same possible worlds. From that perspective, (1) may be more motivated. After all, if P entails Q, then the thought that P is necessarily equivalent to the thought that P and Q. So if one knows P and P entails Q, one automatically knows that P and Q since the thought that P and Q just is, from this perspective, the thought that P. Assuming that knowledge distributes over conjunction – that is, that if one knows P and Q, one knows P and knows Q – then (1) is secured. But, despite its distinguished list of adherents, such a conception of belief is a deviant one, and I shall not pursue it here.

- 2 See, for example, Dretske (1971), Stine (1976), Nozick (1981), Cohen (1988), and Richard Feldman (1999).
- 3 Note that this principle allows for the possibility of some pair P, Q such that P entailed Q but that it was impossible to know P without already knowing Q. Thus the principle at least allows that there be a knowable P and an entailed Q such that deductive inference from P to Q was not a possible route to knowing Q.
- 4 There are other, less decisive, worries as well. Suppose I know P and know that necessarily $P \supset I$ will never perform a deductive inference again. Is it intuitive to say still that I have all it takes, evidentially speaking, to know the consequent?
- 5 For helpful discussion of the relevant issues, see Kripke (1979) and Soames (2000).
- 6 This is a (slightly improved) version of Williamson (2000, p. 117), restricted to single premise inferences. Williamson has an insightful take on the root of epistemic closure intuitions, namely the idea that “deduction is a way of extending one’s knowledge,” one that I take on board in what follows. The notion of “competent deduction” could obviously do with some elaboration, though I won’t attempt that here. Another general issue worth considering (raised in correspondence by Jonathan Kvanig) is whether, in a case where Q is competently deduced from P, a (misleading) defeater for Q – that destroys knowledge of Q – is automatically a defeater for knowledge of P.
- 7 This may be a little quick: one might worry that an exercise of one’s deductive capacities that in fact is competent carries a risk of being incompetent. I shall not address that worry here.
- 8 Such a conception is, for example, defended by Williamson (2000), who explicitly adduces it in defense of generalized closure.
- 9 I leave out the “retaining all along” clause for ease of exposition.
- 10 A similar line of thought shows that Dretske is committed to denying a seemingly humble, restricted version of closure, namely addition closure – that if one knows P then by deduction one can know P or Q. See my *Knowledge and Lotteries* (forthcoming). Nozick (1981) is willing to abandon distribution along with closure. I await Dretske’s opinion on the matter.
- 11 In Nozick’s own case (1981), the “fix” to these anticipated worries famously took the form of an appeal to methods: one knows P, only if, were P false, one wouldn’t have believed P by the same method. It is not my purpose to immerse myself in the details of Nozick’s account here, though many of the remarks that follow are pertinent.
- 12 The main ideas are laid out in Dretske (1971).
- 13 Of course, we epistemologists are often a little lax in our closure-illustrating examples. We may say that the proposition that I have a meeting tomorrow strictly entails that I will not die beforehand, ignoring the possibility that I reincarnate, or go through with the meeting as a ghost. We may say that the proposition that I am seeing a zebra in the cage entails that I am not seeing a cleverly disguised mule, ignoring the possibility that the cage is itself a cleverly disguised mule or else that the mule is cleverly disguised by being inserted within a zebra that is somehow kept alive. So most illustrations of deductive closure aren’t really illustrations of deductive closure. Let us trust that such sloppiness is neither here nor there.
- 14 Best to avoid the pitfalls of devising a necessary and sufficient condition here.
- 15 I am assuming that claims of the form “P wouldn’t be the case unless Q were the case” entail “If Q weren’t the case P wouldn’t be the case.” Dretske certainly relies on such an entailment in generating anti-closure results from his formulation of a conclusive reasons requirement. For the record, I am open to being convinced that “unless” counterfactuals have special features that make trouble for that purported entailment, though I doubt that this would bolster Dretske’s cause.

- 16 Obviously, there are those who will insist on an externalist notion of reasons according to which, say, the fact that there is a zebra may count as a reason (when I see that there is a zebra). If one opted for this notion of reasons, then I may after all have conclusive reasons for thinking that I am not seeing a cleverly disguised mule. (And, *pace* Dretske, note 5, one's reasons would be logically conclusive.) I shall not pursue the matter further here.
- 17 Note that this simulates some of the work of Nozick's own appeal to methods.
- 18 I could run the case with pints of beer, which may be more easily grasped (as long as it is agreed that the world where I drink enough beer to hallucinate that I have drunk a normal amount is a distant one!).
- 19 Note that Dretske is happy enough to use his conclusive reasons framework to explain why I do not know that I will win the lottery (see Dretske, 1971, and "The case against closure," note 7). But it may well be that for any belief about the future, there are outcomes with a real – though low – probability that, if actualized, will render the belief false. The small chance of winning the lottery means, for Dretske, that we do not know we will lose. But what of the small chance postulated by quantum mechanics of the cup flying sideways if I drop it? And what of the small chance that I will have a fatal heart attack? We are in need of an explanation as to how the cases are supposed to be different, unless skepticism about the future is to be embraced. I discuss these issues at length in *Knowledge and Lotteries* and "Chance and counterfactuals" (forthcoming).
- 20 A natural application of Bayesian confirmation theory supports this latter contention. Supposing one's conditional probability of Misprint on the newspaper saying that Manchester United beat Coventry City is N (where N is presumably less than 1), then since one's conditional probability of Misprint on the newspaper denying Manchester United beat Coventry City is 1, one's new rational credence on the evidence that the newspaper said Manchester United beat Coventry City will drop from somewhere between N and 1 to N . (One's new rational credence will be obtained by conditionalization. One's old rational credence will be one's rational expectation that the newspaper will deny that Manchester United won plus N multiplied by one's rational credence that the newspaper will say that Manchester United won, assuming (somewhat unrealistically) that it is certain that the newspaper will deliver a verdict on the matter. If not, adjust "deny" to "not say.")
- 21 Perhaps instead, while one knows and believes one knows, this kind of situation is one where certain beliefs, albeit temporarily, no longer control one's behavior. There are no doubt fertile connections here between epistemology and the relevant issues in philosophical psychology.
- 22 Of course, if one thought knowledge required a kind of conviction over and above "belief" in the ordinary sense, then one would slightly reformulate (4) so that it only covers cases where deductive inference produces the suitable kind of conviction.
- 23 One way to resist is to insist that one's body of evidence is misdescribed by such arguments. That, I take it, is the strategy suggested by Williamson (2000). Perhaps, when I see that there is a table, my evidence is not merely that I seem to see a table, but that there is a table. I shall not explore the matter further here.
- 24 Thanks to Fred Dretske, Peter Klein, Jonathan Kvanig, David Manley, Ernest Sosa, Timothy Williamson, and especially Tamar Gendler for helpful comments and conversation. The ideas in this essay overlap in places with material in my *Knowledge and Lotteries* (forthcoming).

References

- Carroll, L. (1895) What the Tortoise said to Achilles. *Mind*, 4, 278–80.
- Cohen, S. (1988) How to be a fallibilist. *Philosophical Perspectives*, 2: 91–123.
- DeRose, K. (1996). Knowledge, assertion and lotteries. *Australasian Journal of Philosophy*, 84, 568–80.
- Dretske, F. (1970) Epistemic operators. *Journal of Philosophy*, 67, 1007–23.
- Dretske, F. (1971) Conclusive reasons. *Australasian Journal of Philosophy*, 49, 1–22.
- Feldman, R. (1999) Contextualism and skepticism. *Philosophical Perspectives*, 13, 91–114.
- Hawthorne, J. (forthcoming) Chance and Counterfactuals. *Philosophy and Phenomenological Research*.
- Hawthorne, J. (forthcoming) *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Kahneman, D., Slovic, P. and Tversky, A. (1982) *Judgment under Uncertainty: Heuristics and Biases*. Cambridge: Cambridge University Press.
- Kahneman, D. and Tversky, A. (2000) *Choices, Values and Frames*. Cambridge: Cambridge University Press.
- Kripke, S. (1979) A puzzle about belief. In A. Margalit (ed.), *Meaning and Use*. Dordrecht: Reidel.
- Lewis, D. (1973) *Counterfactuals*. Cambridge, MA: Harvard University Press.
- Nozick, R. (1981) *Philosophical Explanations*. Oxford: Oxford University Press.
- Soames, S. (2002) *Beyond Rigidity*. Oxford: Oxford University Press.
- Stine, G. C. (1976) Skepticism, relevant alternatives, and deductive closure. *Philosophical Studies*, 29, 249–61.
- Unger, P. (1975) *Ignorance: A Case for Skepticism*. Oxford: Oxford University Press.
- Williamson, T. (2000) *Knowledge and Its Limits*. Oxford: Oxford University Press.

Reply to Hawthorne

Fred Dretske

John Hawthorne does an admirable job of describing the “costs” connected with my rejection of closure. These costs are, I admit, significant. I would not be willing to pay this price if I thought there were alternatives that were less expensive. Philosophy is a business where one learns to live with spindly brown grass in one’s own yard because neighboring yards are in even worse shape. From where I sit, the grass on the other side of the fence looks worse than mine.

One option, of course, is skepticism. Though some may be willing to embrace it, John and I agree that this is a last resort. If the *only* choices are skepticism or the rejection of closure, then – too bad – we had better learn to live without closure (distribution, etc.).

As a historical footnote, I wasn’t led to deny closure because it represented a way around skepticism. I was led to it because it was a result of what I took to be a plausible condition on the evidence (justification, reasons) required for knowledge. If your reasons for believing P are such that you *might* have them when P is false, then they

aren't good enough to *know* that P is true. You need something more. That is why you can't know you are going to lose a lottery just because your chances of losing are 99.99 percent. Even with those odds, you still might win (someone with those odds against him *will* win). That is why you can't learn – can't come to know – that P is true if all you have to go on is the word of a person who might lie about whether or not P is so. This is just another way of saying that knowledge requires reasons or evidence (in this case, testimony) you wouldn't have if what you end up believing were false. You can learn things from people, yes, but only from people who wouldn't say it unless it were true.

John (in section 3) does a beautiful job in showing the problems with this obvious (to me) condition on knowledge. I had grappled with some of these problems myself, but others, I confess (even after 30 years of defending the view) were completely new to me. Contrary to what he suggests (case 2 in category 2, section 3), though, I think the counterfactual condition accords with common and widespread intuitions about knowledge of the future. In my more careful moments, I prefer to say (on Monday) not that I know that I will be in my office on Wednesday, but that I know that is where I intend or plan to be on Wednesday. I retreat to this more cautious claim because I realize that my reasons for thinking I will be in my office on Wednesday are reasons I would have on Monday even if a Tuesday accident were to prevent me from being there on Wednesday. As long as I might not make it to the office, I do not know I will be there. How could I? One doesn't, generally speaking, have conclusive reasons for believing that something will happen. Death may be an exception, but I doubt whether taxes are. I'm a bit of a skeptic about the future.

But John is certainly right when he says that these counterfactuals are a slippery and shifty business. So a theory of knowledge that appeals to them is exposed to examples of the embarrassing sort that he describes. This, I confess, is a problem. It is not *always* a problem. Sometimes this shiftiness in the way we understand the counterfactual reflects quite neatly – or so I think – the shiftiness and subtlety in our judgments about exactly what and exactly when someone knows something. Does S, standing at the stove watching water boil, know that the water is boiling? Can he see that it is boiling? Well, would it look that way if it weren't boiling? If it wouldn't, he can see that the water is boiling. He can know this despite not being able to see – perhaps not even knowing (but just correctly assuming) – that it is water (not gin) he sees in the pot. The relevant counterfactual is: would the water (not just anything) look that way if it weren't boiling. If *this* counterfactual is true, one can know the water is boiling without knowing (at least not by seeing) that it is water that is boiling since, given the look of things, it could be gin (vodka, pear juice, etc.). I think this example is related in interesting ways to John's lifeguard example, but I won't take the time to spell out the details.

So sometimes the counterfactual condition captures important distinctions. But not always. As John points out, we sometimes seem to have conclusive reasons for believing “manifestly” heavyweight propositions and lack them for clearly lightweight propositions. So things don't sort out as neatly as I would like.

But let's look at the neighbors' yards. How green is their grass? John does not undertake to propose or to defend any particular solution to the problems for which

the denial of closure represents an answer. Without endorsing it, however, he suggests that despite my low opinion (and “flawed understanding”) of contextualism, this approach to understanding knowledge offers some promise in this regard. We can keep closure and avoid skepticism if we think of knowledge as relative to attributional context. To say that it is relative to attributional context is to say that whether S knows something – that she has two hands, for instance – depends on the context of the person who is saying S knows it. If I, a philosopher, worried about brains in vats and Cartesian demons, say it, then S doesn’t know she has two hands. At least what I say when I say she doesn’t know it is true. Why? Because, assuming S is in no position to know she isn’t being deceived by a Cartesian demon, then, given closure, neither can she know that she has two hands. If she knew she had two hands, she would be able to infer, and thereby come to know, by closure, that she isn’t being deceived by such a demon into falsely believing she has two hands. So when I say she doesn’t know she has two hands, what I say is true. Neither do I know I have two hands. Nobody knows it. Skepticism is true. At least it is true when I assert it. But if S, an ordinary person on the street, someone without the least tincture of philosophy, says she knows that she has two hands, what she says is true. She attributes knowledge to herself in an ordinary, practical, context in which demons and handless-brains-in-vats are not relevant possibilities. They aren’t alternatives she means to be excluding, evidentially, when she attributes knowledge to herself. So what she says is true, and it is true, mind you, even when she attributes this knowledge to me, the philosopher pondering skepticism in the seminar room. Whether I know that I have two hands depends not on what alternatives I (the attributee) take to be relevant, but on what possibilities S, the attributor, takes to be relevant. So, if we imagine S (with no tincture of philosophy) asserting that skepticism is false or (given that she probably doesn’t understand what skepticism is) that she and most other people know a great many things about the world, she speaks truly. Skepticism is false when S is the speaker.

So who, according to contextualism, is right? Am I, a philosopher, right when I (given my context) say that nobody knows they have hands. Or is S right? We are, I’m afraid, both right. And that is where my low opinion of contextualism comes from. The only people who can truly assert that skepticism is false are those, like S, who don’t understand what skepticism is or those who understand it but who steadfastly ignore as irrelevant (for purposes of attributing knowledge to themselves and others) exactly those possibilities the skeptic insists are relevant.

I don’t really consider this to be an answer to skepticism or, if it is, it is merely a backhanded way of denying closure. It looks to me as though contextualism succeeds in keeping closure (I use John’s improved statement of closure, i.e. (4)) only by denying it a role in reaching conclusions about the heavyweight propositions the skeptic is concerned with. If you actually try to use closure to conclude from the fact that you have two hands that you are not being massively deceived, then you automatically change the context in which you attribute knowledge to yourself (the context is now one in which the possibility of being deceived in this undetectable way is being explicitly raised) and, given closure (and the fact that you can’t know you aren’t being deceived in this undetectable way), you cease to know the premise (that you have two hands) you used to reach that heavyweight conclusion. So, despite reaching this

heavyweight conclusion in a perfectly respectable way from a premise you (when you started) knew to be true, you do not wind up knowing that the conclusion is true – that you aren't being deceived. Instead, you lose the knowledge of the premise (that you have two hands) you used to reach that conclusion. This grass doesn't look very green to me.