

Journal of Philosophy, Inc.

Epistemic Operators

Author(s): Fred I. Dretske

Reviewed work(s):

Source: *The Journal of Philosophy*, Vol. 67, No. 24 (Dec. 24, 1970), pp. 1007-1023

Published by: [Journal of Philosophy, Inc.](#)

Stable URL: <http://www.jstor.org/stable/2024710>

Accessed: 13/09/2012 20:26

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at
<http://www.jstor.org/page/info/about/policies/terms.jsp>

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



Journal of Philosophy, Inc. is collaborating with JSTOR to digitize, preserve and extend access to *The Journal of Philosophy*.

<http://www.jstor.org>

THE JOURNAL OF PHILOSOPHY

VOLUME LXVII, NO. 24, DECEMBER 24, 1970

EPISTEMIC OPERATORS *

SUPPOSE Q is a necessary consequence of P . Given only this much, it is, of course, quite trivial that if it is true that P , then it must also be true that Q . If it is a fact that P , then it must also be a fact that Q . If it is necessary that P , then it is necessary that Q ; and if it is possible that P , then it must also be possible that Q .

I have just mentioned four prefixes: 'it is true that', 'it is a fact that', 'it is necessary that', and 'it is possible that'. In this paper I shall refer to such affixes as *sentential operators* or simply *operators*; when affixed to a sentence or statement, they operate on it to generate another sentence or statement. The distinctive thing about the four operators I have just mentioned is that, if Q is a necessary consequence of P , then the statement we get by operating on Q with one of these four operators is a necessary consequence of the statement we get by operating on P with the same operator. This may be put more succinctly if we let 'O' stand for the operator in question and 'O(P)' for the statement we get by affixing the operator 'O' to the statement ' P '. We can now say that the above four operators share the following property: if P entails Q , then O(P) entails O(Q). I shall call any operator having this property a *penetrating operator* (or, when emphasis is required, a *fully penetrating operator*). In operating on P these operators penetrate to every necessary consequence of P .

We are now in a position to ask ourselves a preliminary question. The answer to this question is easy enough, but it will set the

* Versions of this paper were read to the philosophy departments of several universities in the United States and Canada during the year 1969/70. I profited greatly from these discussions. I wish especially to thank Paul Dietl who helped me to see a number of points more clearly (perhaps still not clearly enough in his opinion). Finally, my exchanges with Mr. Don Affeldt were extremely useful; I am much indebted to him in connection with some of the points made in the latter portions of the paper.

stage for more difficult questions. Are all sentential operators fully penetrating operators? Are all operators such that if P entails Q , then $O(P)$ entails $O(Q)$? If *all* operators are penetrating operators, then each of the following statements must be true (when P entails Q):

- (1) You cannot have a reason to believe that P unless you have a reason to believe that Q .
- (2) You cannot know that P unless you know that Q .
- (3) You cannot explain why P is the case unless you can explain why Q is the case.
- (4) If you assert that P , then you assert that Q .
- (5) If you hope that P , then you hope that Q .
- (6) If it is strange (or accidental) that P , then it must be strange (or accidental) that Q .
- (7) If it was a mistake that P , then it was a mistake that Q .

This list begins with two epistemic operators, 'reason to believe that' and 'know that'. Since I shall be concerned with these later in the paper, let me skip over them now and look at those appearing near the end of the list. They will suffice to answer our opening question, and their status is much less problematic than that of some of the other operators.

'She lost' entails 'Someone lost'. Yet, it may be strange that she lost, not at all strange that someone lost. 'Bill and Susan married each other' entails that Susan got married; yet, it may be quite odd that (strange that, incredible that) Bill and Susan married each other but quite unremarkable, not at all odd that, Susan got married. It may have been a mistake that they married each other, not a mistake that Susan got married. Or finally, 'I hit the bull's-eye' entails that I either hit the bull's-eye or the side of the barn; and though I admit that it was lucky that (accidental that) I hit the bull's-eye, I will deny that it was lucky, an accident, that I hit either the bull's-eye or the side of the barn.

Such examples show that not all operators are fully penetrating. Indeed, such operators as 'it is strange that', 'it is accidental that' and 'it is a mistake that' fail to penetrate to some of the most elementary logical consequences of a proposition. Consider the entailment between ' $P \cdot Q$ ' and ' Q '. Clearly, it may be strange that P and Q , not at all strange that P , and not at all strange that Q . A concatenation of factors, no one of which is strange or accidental, may itself be strange or accidental. Taken by itself, there is nothing odd or suspicious about Frank's holding a winning ticket in the first race. The same could be said about any of the other races: there

is nothing odd or suspicious about Frank's holding a winning ticket in the n th race. Nonetheless, there is something very odd, very suspicious, in Frank's having a winning ticket in n races.

Therefore, not only are these operators *not* fully penetrating, they lie, as it were, on the other end of the spectrum. They fail to penetrate to some of the most elementary consequences of a proposition. I shall refer to this class of operators as *nonpenetrating* operators. I do not wish to suggest by this label that such operators are totally impotent in this respect (or that they are all uniform in their degree of penetration). I mean it, rather, in a rough, comparative, sense: their *degree of penetration* is less than that of any of the other operators I shall have occasion to discuss.

We have, then, two ends of the spectrum with examples from both ends. Anything that falls between these two extremes I shall call a *semi-penetrating operator*. And with this definition I am, finally, in a position to express my main point, the point I wish to defend in the rest of this paper. It is, simply, that all epistemic operators are semi-penetrating operators. There is both a trivial and a significant side to this claim. Let me first deal briefly with the trivial aspect.

The epistemic operators I mean to be speaking about when I say that all epistemic operators are semi-penetrating include the following:

- (a) S knows that . . .
- (b) S sees (or can see) that . . .
- (c) S has reason (or a reason) to believe that . . .
- (d) There is evidence to suggest that . . .
- (e) S can prove that . . .
- (f) S learned (discovered, found out) that . . .
- (g) In relation to our evidence it is probable that . . .

Part of what needs to be established in showing that these are all semi-penetrating operators is that they all possess a degree of penetration greater than that of the nonpenetrating operators. This is the trivial side of my thesis. I say it is trivial because it seems to me fairly obvious that if someone knows that P and Q , has a reason to believe that P and Q , or can prove that P and Q , he thereby knows that Q , has a reason to believe that Q , or can prove (in the appropriate epistemic sense of this term) that Q . Similarly, if S knows that Bill and Susan married each other, he (must) know that Susan got married (married someone). If he knows that P is the case, he knows that P or Q is the case (where the 'or' is understood in a sense which makes ' P or Q ' a necessary consequence of ' P '). This is not a claim

about what it would be appropriate to say, what the person himself thinks he knows or would say he knows. It is a question, simply, of what he knows. It may not be appropriate to say to Jim's wife that you know it was either her husband, Jim, or Harold who sent the neighbor lady an expensive gift *when you know it was Harold*. For, although you do know this, it is misleading to say you know it—especially to Jim's wife.

Let me accept, therefore, without further argument that the epistemic operators are not, unlike 'lucky that', 'strange that', 'a mistake that', and 'accidental that', nonpenetrating operators. I would like to turn, then, to the more significant side of my thesis. Before I do, however, I must make one point clear lest it convert my entire thesis into something as trivial as the first half of it. When we are dealing with the epistemic operators, it becomes crucial to specify whether the agent in question knows that *P* entails *Q*. That is to say, *P* may entail *Q*, and *S* may know that *P*, but he may not know that *Q* *because*, and perhaps *only* because, he fails to appreciate the fact that *P* entails *Q*. When *Q* is a simple logical consequence of *P* we do not expect this to happen, but when the propositions become very complex, or the relationship between them very complex, this might easily occur. Let *P* be a set of axioms, *Q* a theorem. *S*'s knowing *P* does not entail *S*'s knowing *Q* just because *P* entails *Q*; for, of course, *S* may not know that *P* entails *Q*, may not know that *Q* is a *theorem*. Hence, our epistemic operators will turn out *not* to be penetrating because, and perhaps *only* because, the agents in question are not fully cognizant of all the implications of what they know to be the case, can see to be the case, have a reason to believe is the case, and so on. Were we all ideally astute logicians, were we all fully apprised of all the necessary consequences (supposing this to be a well defined class) of every proposition, perhaps then the epistemic operators would turn into fully penetrating operators. That is, assuming that if *P* entails *Q*, we *know* that *P* entails *Q*, then every epistemic operator is a penetrating operator: the epistemic operators penetrate to all the *known* consequences of a proposition.

It is this latter, slightly modified, claim that I mean to reject. Therefore, I shall assume throughout the discussion that when *Q* is a necessary consequence of *P*, every relevant agent *knows that it is*. I shall be dealing with only the *known consequences* (in most cases because they are immediate and obvious consequences). What I wish to show is that, even under this special restriction, the epistemic operators are *only* semi-penetrating.

I think many philosophers would disagree with this contention. The conviction is that the epistemic worth of a proposition is hereditary under entailment, that whatever the epistemic worth of *P*, *at least* the same value must be accorded the known consequences of *P*. This conviction finds expression in a variety of ways. Epistemic logic: if *S* knows that *P*, and knows that *P* entails *Q*, then *S* knows that *Q*. Probability theory: if *A* is probable, and *B* is a logical consequence of *A*, then *B* is probable (relative to the same evidence, of course). Confirmation theory: if evidence *e* tends to confirm hypothesis *h*, then *e* indirectly confirms all the logical consequences of *h*. But perhaps the best evidence in favor of supposing that most philosophers have taken the epistemic operators to be fully penetrating is the way they have argued and the obvious assumptions that structure their arguments. Anyone who has argued in the following way seems to me to be assuming the thesis of penetrability (as I shall call it): if you do not know whether *Q* is true or not, and *P* cannot be true unless *Q* is true, then you (obviously) do not know whether *P* is true or not. A slightly more elaborate form of the same argument goes like this: If *S* does not know whether or not *Q* is true, then for all he knows it might be false. If *Q* is false, however, then *P* must also be false. Hence, for all *S* knows, *P* may be false. Therefore, *S* does not know that *P* is true. This pattern of argument is sprinkled throughout the epistemological literature. Almost all skeptical objections trade on it. *S* claims to know that this is a tomato. A necessary consequence of its being a tomato is that it is not a clever imitation which only looks and feels (and, if you will, tastes) like a tomato. But *S* does not know that it is *not* a clever imitation that only looks and feels (and tastes) like a tomato. (I assume here that no one is prepared to argue that anything that looks, feels, and tastes like a tomato to *S* *must be* a tomato.) Therefore, *S* does not know that this is a tomato. We can, of course, reply with G. E. Moore that we certainly *do* know it is a tomato (after such an examination) and since tomatoes are not imitations we know that this is not an imitation. It is interesting to note that this reply presupposes the same principle as does the skeptical objection: they both assume that if *S* knows that this is a *P*, and knows that every *P* is a *Q*, then *S* knows that this is a *Q*. The only difference is that the skeptic performs a modus tollens, Moore a modus ponens. Neither questions the principle itself.

Whether it be a question of dreams or demons, illusions or fakes, the same pattern of argument emerges. If you know this is a chair,

you must know that you are not dreaming (or being deceived by a cunning demon), since its being a (real) chair entails that it is not simply a figment of your own imagination. Such arguments assume that the epistemic operators, and in particular the operator 'to know', penetrate to all the known consequences of a proposition. If these operators were not penetrating, many of these objections might be irrelevant. Consider the following exchange:

S: How strange! There are tomatoes growing in my apple tree.

K: That isn't strange at all. Tomatoes, after all, are physical objects and what is so strange about physical objects growing in your apple tree?

What makes K's reply so silly is that he is treating the operator 'strange that' as a fully penetrating operator: it cannot be strange that there are tomatoes growing in the apple tree unless the consequences of this (e.g., there are objects growing in your apple tree) are also strange. Similarly, it *may not* be at all relevant to object to someone who claims to know that there are tomatoes in the apple tree that he does not know, cannot be absolutely sure, that there are really any material objects. Whether or not this is a relevant objection will depend on whether or not this particular consequence of there being tomatoes in the apple tree is one of the consequences to which the epistemic operators penetrate. What I wish to argue in the remainder of this paper is that the traditional skeptical arguments exploit precisely those consequences of a proposition to which the epistemic operators do not penetrate, precisely those consequences which distinguish the epistemic operators from the fully penetrating operators.

In support of this claim let me begin with some examples which are, I think, fairly intuitive and then turn to some more problematic cases. I shall begin with the operator 'reason to believe that' although what I have to say could be said as well with any of them. This particular operator has the added advantage that if it can be shown to be only semi-penetrating, then many accounts of knowledge, those which interpret it as a form of justified true belief, would also be committed to treating 'knowing that' as a semi-penetrating operator. For, presumably, 'knowing that' would not penetrate any deeper than one's 'reasons for believing that'.

Suppose you have a reason to believe that the church is empty. *Must* you have a reason to believe that it is a church? I am not asking whether you generally have such a reason. I am asking whether one can have a reason to believe the church empty without having a reason to believe that it is a church which is empty. Certainly your

reason for believing that the church is empty is not *itself* a reason to believe it is a church; or it *need not* be. Your reason for believing the church to be empty may be that you just made a thorough inspection of it without finding anyone. That is a good reason to believe the church empty. Just as clearly, however, it is not a reason, much less a good reason, to believe that what is empty is a church. The fact is, or so it seems to me, I do not have to have *any* reason to believe it is a church. Of course, I would never *say* the church was empty, or that I had a reason to believe that the church was empty, unless I believed, and presumably had a reason for so believing, that *it was* a church which was empty, but this is a presumed condition of my *saying* something, not of my having a reason to believe something. Suppose I had simply assumed (correctly as it turns out) that the building was a church. Would this show that I had no reason to believe that the church was empty?

Suppose I am describing to you the "adventures" of my brother Harold. Harold is visiting New York for the first time, and he decides to take a bus tour. He boards a crowded bus and immediately takes the last remaining seat. The little old lady he shouldered aside in reaching his seat stands over him glowering. Minutes pass. Finally, realizing that my brother is not going to move, she sighs and moves resignedly to the back of the bus. Not much of an adventure, but enough, I hope, to make my point. I said that the little old lady realized that my brother would not move. Does this imply that she realized that, or knew that, *it was my brother* who refused to move? Clearly not. We can say that *S* knows that *X* is *Y* without implying that *S* knows that *it is X* which is *Y*. We do not *have* to describe our little old lady as knowing that *the man* or *the person* would not move. We can say that she realized that, or knew that, *my brother* would not move (minus, of course, this pattern of emphasis), and we can say this because saying this does not entail that the little old lady knew that, or realized that, it was my brother who refused to move. She knew that my brother would not move, and she knew this despite the fact that she did not know something that was necessarily implied by what she did know—viz., that the person who refused to move was my brother.

I have argued elsewhere that to see that *A* is *B*, that the roses are wilted for example, is not to see, not even to be able to see, that they are roses which are wilted.¹ To see that the widow is limping is not to see that it is a widow who is limping. I am now arguing

¹ *Seeing and Knowing* (Chicago: University Press, 1969), pp. 93–112, and also "Reasons and Consequences," *Analysis* (April 1968).

that this same feature holds for all epistemic operators. I can know that the roses are wilting without knowing that they are roses, know that the water is boiling without knowing that it is water, and prove that the square root of 2 is smaller than the square root of 3 and, yet, be unable to prove what is entailed by this—viz., that the number 2 *has* a square root.

The general point may be put this way: there are certain presuppositions associated with a statement. These presuppositions, although their truth is entailed by the truth of the statement, are not part of what is *operated on* when we operate on the statement with one of our epistemic operators. The epistemic operators do not *penetrate to* these presuppositions. For example, in saying that the coffee is boiling I assert that the coffee is boiling, but in asserting this I do not assert that *it is* coffee which is boiling. Rather, this is taken for granted, assumed, presupposed, or what have you. Hence, when I say that I have a reason to believe that the coffee is boiling, I am not saying that this reason applies to the fact that it is coffee which is boiling. This is *still* presupposed. I may have such a reason, of course, and chances are good that I do have such a reason or I would not have referred to what I believe to be boiling *as coffee*, but to have a reason to believe the coffee is boiling is not, thereby, to have a reason to believe it is coffee which is boiling.

One would expect that if this is true of the semi-penetrating operators, then it should also be true of the nonpenetrating operators. They also should fail to reach the presuppositions. This is exactly what we find. It may be accidental that the two trucks collided, but not at all accidental that it was two trucks that collided. Trucks were the only vehicles allowed on the road that day, and so it was not at all accidental or a matter of chance that the accident took place between two trucks. Still, it was an accident that the two trucks collided. Or suppose Mrs. Murphy mistakenly gives her cat some dog food. It need not be a mistake that she gave the food to *her* cat, or *some* food to *a* cat. This was intentional. What was a mistake was that it was dog food that she gave to her cat.

Hence, the first class of consequences that differentiate the epistemic operators from the fully penetrating operators is the class of consequences associated with the presuppositions of a proposition. The fact that the epistemic operators do not penetrate to these presuppositions is what helps to make them semi-penetrating. And this is an extremely important fact. For it would appear that if this is true, then to know that the flowers are wilted I do not have to know that they are flowers (which are wilted) and, therefore, do not have

to know all those consequences which follow from the fact that they are flowers, real flowers, which I know to be wilted.

Rather than pursue this line, however, I would like to turn to what I consider to be a more significant set of consequences—"more significant" because they are the consequences that are directly involved in most skeptical arguments. Suppose we assert that x is A . Consider some predicate, ' B ', which is incompatible with A , such that nothing can be both A and B . It then follows from the fact that x is A that x is not B . Furthermore, if we conjoin B with any other predicate, Q , it follows from the fact that x is A that x is not- $(B$ and $Q)$. I shall call this type of consequence a *contrast consequence*, and I am interested in a particular subset of these; for I believe the most telling skeptical objections to our ordinary knowledge claims exploit a particular set of these contrast consequences. The exploitation proceeds as follows: someone purports to know that x is A , that the wall is red, say. The skeptic now finds a predicate ' B ' that is incompatible with ' A '. In this particular example we may let ' B ' stand for the predicate 'is white'. Since ' x is red' entails ' x is not white' it also entails that x is not- $($ white and $Q)$ where ' Q ' is any predicate we care to select. Therefore, the skeptic selects a ' Q ' that gives expression to a condition or circumstance under which a white wall would appear exactly the same as a red wall. For simplicity we may let ' Q ' stand for: 'cleverly illuminated to look red'. We now have this chain of implications: ' x is red' entails ' x is not white' entails ' x is not white cleverly illuminated to look red'. If 'knowing that' is a penetrating operator, then if anyone knows that the wall is red he must know that it is not white cleverly illuminated to look red. (I assume here that the relevant parties know that if x is red, it cannot be white made to look red.) He must know that this particular contrast consequence is true. The question is: do we, generally speaking, know anything of the sort? Normally we never take the trouble to check the lighting. We seldom acquire any *special* reasons for believing the lighting normal although we can talk vaguely about there being no reason to think it unusual. The fact is that we habitually take such matters for granted, and although we normally have *good* reasons for making such routine assumptions, I do not think these reasons are sufficiently good, not without special precautionary checks in the particular case, to say of the particular situation we are in that we *know* conditions are normal. To illustrate, let me give you another example—a silly one, but no more silly than a great number of skeptical arguments with which we are all familiar. You take your

son to the zoo, see several zebras, and, when questioned by your son, tell him they are zebras. Do you know they are zebras? Well, most of us would have little hesitation in saying that we did know this. We know what zebras look like, and, besides, this is the city zoo and the animals are in a pen clearly marked "Zebras." Yet, something's being a zebra implies that it is not a mule and, in particular, not a mule cleverly disguised by the zoo authorities to look like a zebra. Do you know that these animals are not mules cleverly disguised by the zoo authorities to look like zebras? If you are tempted to say "Yes" to this question, think a moment about what reasons you have, what evidence you can produce in favor of this claim. The evidence you *had* for thinking them zebras has been effectively neutralized, since it does not count toward their *not* being mules cleverly disguised to look like zebras. Have you checked with the zoo authorities? Did you examine the animals closely enough to detect such a fraud? You might do this, of course, but in most cases you do nothing of the kind. You have some general uniformities on which you rely, regularities to which you give expression by such remarks as, "That isn't very likely" or "Why should the zoo authorities do that?" Granted, the hypothesis (if we may call it that) is not very plausible, given what we know about people and zoos. But the question here is not whether this alternative is plausible, not whether it is more or less plausible than that there are real zebras in the pen, but whether *you know* that this alternative hypothesis is false. I don't think you do. In this I agree with the skeptic. I part company with the skeptic only when he concludes from this that, therefore, you do not know that the animals in the pen are zebras. I part with him because I reject the principle he uses in reaching this conclusion—the principle that if you do not know that *Q* is true, when it is known that *P* entails *Q*, then you do not know that *P* is true.

What I am suggesting is that we simply admit that we do *not* know that some of these contrasting "skeptical alternatives" are *not* the case, but refuse to admit that we do not know what we originally said we knew. My knowing that the wall is red certainly entails that the wall is red; it also entails that the wall is not white and, in particular, it entails that the wall is not white cleverly illuminated to look red. But it does not follow from the fact that I know that the wall is red that I *know* that it is not white cleverly illuminated to look red. Nor does it follow from the fact that I know that those animals are zebras that I know that they are not mules cleverly disguised to look like zebras. These are some of the

contrast consequences to which the epistemic operators do not penetrate.

Aside from asserting this, what arguments can be produced to support it? I could proceed by multiplying examples, but I do not think that examples alone will support the full weight of this view. The thesis itself is sufficiently counterintuitive to render controversial most of the crucial examples. Anyone who is already convinced that skepticism is wrong and who is yet troubled by the sorts of skeptical arguments I have mentioned will, no doubt, take this itself as an argument in favor of my claim that the epistemic operators are only semi-penetrating. This, however, hardly constitutes an argument against skepticism. For this we need *independent* grounds for thinking that the epistemic operators do not penetrate to the contrast consequences. So I shall proceed in a more systematic manner. I shall offer an analogy with three other operators and conclude by making some general remarks about what I think can be learned from this analogy. The first operator is 'explains why' or, more suggestively (for the purposes of this analogy):

(A) *R* is the reason (explanatory reason) that (or why) . . .

For example, the reason why *S* quit smoking was that he was afraid of getting cancer. The second operator has to do with reasons again, but in this case it is a reason which tends to *justify* one in doing something:

(B) *R* is a reason for . . . (*S* to do *Y*).²

For example, the fact that they are selling the very same (type of) car here much more cheaply than elsewhere is a reason to buy it here rather than elsewhere. The status of this as a reason will, of course, depend on a variety of circumstances, but situations can easily be imagined in which this would be a reason for someone to buy the car here. Finally, there is a particular modal relationship which may be construed as a sentential operator:

(C) *R* would not be the case unless . . .

For example, he would not have bid seven no-trump unless he had all four aces. I shall abbreviate this operator as ' $R \rightarrow . . .$ '; hence, our example could be written 'he bid seven no-trump \rightarrow he had all four aces'.

² Unlike our other operators, this one does not have a propositional operand. Despite the rather obvious differences between this case and the others, I still think it useful to call attention to its analogous features.

Each of these operators has features similar to those of our epistemic operators. If one retraces the ground we have already covered, one will find, I think, that these operators all penetrate deeper than the typical nonpenetrating operator. If *R* explains why (or is the reason that) *P* and *Q* are the case, then it explains why (is the reason that) *Q* is the case.³ If I can explain why Bill and Harold are always invited to every party, I can explain why Harold is always invited to every party. From the fact that it was a mistake for me to quit my job it does not follow that it was a mistake for me to do something, but if I had a reason to quit my job, it does follow that I had a reason to do something. And if the grass would not be green unless it had plenty of sunshine and water, it follows that it would not be green unless it had water.

Furthermore, the similarities persist when one considers the presuppositional consequences. I argued that the epistemic operators fail to penetrate to the presuppositions; the above three operators display the same feature. In explaining why he takes his lunch to work, I do not (or need not) explain why he goes to work or why he works at all. The explanation may be obvious in some cases, of course, but the fact is I need not be able to explain why he works (he is *so* wealthy) to explain why he takes his lunch to work (the cafeteria food is *so* bad). The reason why the elms on Main Street are dying is *not* the reason there are elms on Main Street. I have a reason to feed my cat, no reason (not, at least, the same reason) to have a cat. And although it is quite true that he would not have known about our plans if the secretary had not told him, it does not follow that he would not have known about our plans if *some-one other than the secretary* had told him. That is, (He knew about our plans) \rightarrow (The secretary told him) even though it is *not* true that (He knew about our plans) \rightarrow (It was the secretary who told him). Yet, the fact that *it was the secretary* who told him is (I take it) a presuppositional consequence of the fact that *the secretary* told him. Similarly, if George is out to set fire to the first empty building he finds, it may be true to say that George would not have set fire to the church unless it (the church) was empty, yet false to say

³ One must be careful not to confuse sentential conjunction with similar-sounding expressions involving a relationship between two things. For example, to say Bill and Susan got married (if it is intended to mean that they married *each other*), although it entails that Susan got married, does not do so by *simplification*. 'Reason why' penetrates through logical simplification, *not* through the type of entailment represented by these two propositions. That is, the reason they got married is that they loved each other; that they loved each other is not the reason Susan got married.

that George would not have set fire to the church unless *it was a church*.

I now wish to argue that these three operators do not penetrate to a certain set of contrast consequences. To the extent that the epistemic operators are similar to these operators, we may then infer, by analogy, that they also fail to penetrate to certain contrast consequences. This is, admittedly, a weak form of argument, depending as it does on the grounds there are for thinking that the above three operators and the epistemic operators share the same logic in this respect. Nonetheless, the analogy is revealing. Some may even find it persuasive.⁴

(A) The pink walls in my living room clash with my old green couch. Recognizing this, I proceed to paint the walls a compatible shade of green. This is the reason I have, and give, for painting the walls green. Now, in having this explanation for why I painted the walls green, I do not think I have an explanation for two other things, both of which are entailed by what I do have an explanation for. I have not explained why I did not, *instead* of painting the walls green, buy a new couch or cover the old one with a suitable slip cover. Nor have I explained why, instead of painting the walls green, I did not paint them white and illuminate them with green light. The same effect would have been achieved, the same purpose would have been served, albeit at much greater expense.

I expect someone to object as follows: although the explanation given for painting the walls green does not, by itself, explain why the couch was not changed instead, it nonetheless succeeds as an explanation for why the walls were painted green only in so far as there is an explanation for why the couch was not changed instead. If there is no explanation for why I did not change the couch instead, there has been no real, no complete, examination for why the walls were painted green.

I think this objection wrong. I may, of course, have an explanation for why I did not buy a new couch: I love the old one or it has sentimental value. But then again I may not. It just never occurred to me to change the couch; or (if someone thinks that its not oc-

⁴ I think that those who are inclined to give a causal account of knowledge should be particularly interested in the operator ' $R \rightarrow . . .$ ' since, presumably, it will be involved in many instances of knowledge ("many" not "all," since one might wish to except some form of immediate knowledge—knowledge of one's own psychological state—from the causal account). If this operator is only semi-penetrating, then any account of knowledge that relies on the relationship expressed by this operator (as I believe causal accounts must) will be very close to giving a "semi-penetrating" account of 'knowing that'.

curing to me *is* an explanation of why I did not change the couch) I may have thought of it but decided, for what reasons (if any) I cannot remember, to keep the couch and paint the walls. That is to say, I cannot explain why I did not change the couch. I thought of it but I did not do it. I do not know why. Still, I *can* tell you why I painted the walls green. They clashed with the couch.

(B) The fact that they are selling Xs so much more cheaply here than elsewhere may be a reason to buy your Xs here, but it certainly need not be a reason to do what is a necessary consequence of *buying* your Xs here—viz., not *stealing* your Xs here.

(C) Let us suppose that *S* is operating in perfectly normal circumstances, a set of circumstances in which it is true to say that the wall he sees would not (now) look green to him unless it was green (if it were any other color it would look different to him). Although we can easily imagine situations in which this is true, it does not follow that the wall would not (now) look green to *S* if it were white cleverly illuminated to look green. That is,

- (i) The wall looks green (to *S*) \rightarrow the wall is green.
- (ii) The wall is green *entails* the wall is not white cleverly illuminated to look green (to *S*).

are both true; yet, it is *not true* that

- (iii) The wall looks green (to *S*) \rightarrow the wall is not white cleverly illuminated to look green (to *S*).

There are dozens of examples that illustrate the relative impenetrability of this operator. We can truly say that *A* and *B* would not have collided if *B* had not swerved at the last moment and yet concede that they would have collided without any swerve on the part of *B* if the direction in which *A* was moving had been suitably altered in the beginning.⁵

⁵ The explanation for why the modal relationship between *R* and *P* ($R \rightarrow P$) fails to carry over (penetrate) to the logical consequences of *P* (i.e., $R \rightarrow Q$ where *Q* is a logical consequence of *P*) is to be found in the set of circumstances that are taken as *given*, or *held fixed*, in subjunctive conditionals. There are certain logical consequences of *P* which, by bringing in a reference to circumstances tacitly held fixed in the original subjunctive ($R \rightarrow P$), introduce a possible variation in these circumstances and, hence, lead to a *different framework* of fixed conditions under which to assess the truth of $R \rightarrow Q$. For instance, in the last example in the text, when it is said that *A* and *B* would not have collided if *B* had not swerved at the last moment, the truth of this conditional clearly takes it *as given* that *A* and *B* possessed the prior trajectories they in fact had on the occasion in question. *Given* certain facts, including the fact that they were traveling in the direction they were, they would not have collided if *B* had not swerved. Some of the logical consequences of the statement

The structure of these cases is virtually identical with that which appeared in the case of the epistemic operators, and I think by looking just a little more closely at this structure we can learn something very fundamental about our class of epistemic operators and, in particular, about what it means to know something. If I may put it this way, within the context of these operators no fact is an island. If we are simply rehearsing the facts, then we can say that it is a fact that Brenda did not take any dessert (though it was included in the meal). We can say this without a thought about what sort of person Brenda is or what she might have done had she ordered dessert. However, if we put this fact into, say, an explanatory context, if we try to explain this fact, it suddenly appears within a network of related facts, a network of possible alternatives which serve to define *what it is that is being explained*. What is being explained is a function of two things—not only the fact (Brenda did not order any dessert), but also the range of relevant alternatives. A relevant alternative is an alternative that might have been realized in the existing circumstances if the actual state of affairs had not materialized.⁶ When I explain why Brenda did not order any dessert by saying that she was full (was on a diet, did not like anything on the dessert menu), I explain why she did not order any dessert *rather than, as opposed to, or instead of* ordering some dessert and *eating it*. It is this competing possibility which helps to define what it is that I am explaining when I explain why Brenda did not order any dessert. Change this contrast, introduce a different set of relevant alternatives, and you change what it is that is being explained and, therefore, what counts as an explanation, even though (as it were) the same fact is being explained. Consider the following contrasts: ordering some dessert and throwing it at the waiter; ordering some

that *B* swerved do not, however, leave these conditions unaltered—e.g., *B* did not move in a perfectly straight line in a direction 2° counterclockwise to the direction it actually moved. This consequence “tinkers” with the circumstances originally taken *as given* (held fixed), and a failure of penetration will usually arise when this occurs. It *need not be true* that *A* and *B* would not have collided if *B* had moved in a perfectly straight line in a direction 2° counterclockwise to the direction it actually moved.

⁶ I am aware that this characterization of “a relevant alternative” is not, as it stands, very illuminating. I am not sure I can make it more precise. What I am after can be expressed this way: if Brenda *had* ordered dessert, she *would not* have thrown it at the waiter, stuffed it in her shoes, or taken it home to a sick friend (she has no sick friend). These are not alternatives that *might* have been realized in the existing circumstances if the actual state of affairs had not materialized. Hence, they are not relevant alternatives. In other words, the ‘might have been’ in my characterization of a relevant alternative will have to be unpacked in terms of counterfactuals.

dessert and taking it home to a sick friend. With these contrasts none of the above explanations are any longer explanations of why Brenda did not order dessert. Anyone who really wants to know why Brenda did not order dessert and throw it at the waiter will not be helped by being told that she was full or on a diet. This is only to say that, within the context of explanation and within the context of our other operators, the proposition on which we operate must be understood as embedded within a matrix of relevant alternatives. We explain why *P*, but we do so within a framework of competing alternatives *A*, *B*, and *C*. Moreover, if the possibility *D* is not within this contrasting set, not within this network of relevant alternatives, then even though not-*D* follows necessarily from the fact, *P*, which we do explain, we do not explain why not-*D*. Though the fact that Brenda did not order dessert and throw it at the waiter follows necessarily from the fact that she did not order dessert (the fact that is explained), this necessary consequence is not explained by the explanation given. The only contrast consequences to which this operator penetrates are those which figured in the original explanation as relevant alternatives.

So it is with our epistemic operators. To know that *x* is *A* is to know that *x* is *A* within a framework of relevant alternatives, *B*, *C*, and *D*. This set of contrasts, together with the fact that *x* is *A*, serve to define what it is that is known when one knows that *x* is *A*. One cannot change this set of contrasts without changing what a person is said to know when he is said to know that *x* is *A*. We have subtle ways of shifting these contrasts and, hence, changing what a person is said to know *without changing the sentence that we use to express what he knows*. Take the fact that Lefty killed Otto. By changing the emphasis pattern we can invoke a different set of contrasts and, hence, alter what it is that *S* is said to know when he is said to know that Lefty killed Otto. We can say, for instance, that *S* knows that *Lefty* killed Otto. In this case (and I think this is the way we usually hear the sentence when there is no *special* emphasis) we are being told that *S* knows the identity of Otto's killer, that *it was Lefty* who killed Otto. Hence, we expect *S*'s reasons for believing that Lefty killed Otto to consist in facts that single out Lefty as the assailant *rather than* George, Mike, or someone else. On the other hand, we can say that *S* knows that Lefty *killed* Otto. In this case we are being told that *S* knows *what Lefty did to Otto*; he killed him *rather than* merely injuring him, killed him *rather than* merely threatening him, etc. A good reason for believing that Lefty *killed* Otto (rather than merely injuring him) is that Otto is dead,

but this is not much of a reason, if it is a reason at all, for believing that *Lefty* killed Otto. Changing the set of contrasts (from 'Lefty rather than George or Mike' to 'killed rather than injured or threatened') by shifting the emphasis pattern changes what it is that one is alleged to know when one is said to know that Lefty killed Otto.⁷ The same point can be made here as we made in the case of explanation: the operator will penetrate *only* to those contrast consequences which form part of the network of relevant alternatives structuring the original context in which a knowledge claim was advanced. Just as we have not explained why Brenda did not order some dessert and throw it at the waiter when we explained why she did not order some dessert (although what we have explained—her not ordering any dessert—entails this), so also in knowing that Lefty *killed* Otto (knowing that what Lefty did to Otto was kill him) we do not *necessarily* (although we may) know that *Lefty* killed Otto (know that *it was Lefty* who killed Otto). Recall the example of the little old lady who knew that my brother would not move without knowing that it was my brother who would not move.

The conclusions to be drawn are the same as those in the case of explanation. Just as we can say that within the original setting, within the original framework of alternatives that defined what we were trying to explain, we *did explain* why Brenda did not order any dessert, so also within the original setting, within the set of contrasts that defined what it was we were claiming to know, we *did know* that the wall was red and *did know* that it was a zebra in the pen.

To introduce a novel and enlarged set of alternatives, as the skeptic is inclined to do with our epistemic claims, is to exhibit consequences of what we know, or have reason to believe, which we may not know, may not have a reason to believe; but it does not show that we did not know, did not have a reason to believe, whatever it is that has these consequences. To argue in this way is, I submit, as much a mistake as arguing that we have not explained why Brenda did not order dessert (within the original, normal, setting) because we did not explain why she did not order some and throw it at the waiter.

FRED I. DRETSKE

University of Wisconsin

⁷ The same example works nicely with the operator ' $R \rightarrow . . .$ '. It may be true to say that Otto would not be dead unless Lefty *killed* him (unless what Lefty did to him was kill him) without its being true that Otto would not be dead unless *Lefty* killed him (unless it was Lefty who killed him).